

AI Buildout Risk in Private Real Assets *

Gregory Brown
gregwbrown@unc.edu

Christian Lundblad
christian_lundblad@unc.edu

William M. Volckmann II
wmvolckmann@unc.edu

November 20, 2026

Abstract

Institutional investors hold private real assets, such as infrastructure and real estate, as relatively diversified, long-duration, inflation-linked investments. The capital expenditure program supporting artificial intelligence is being built with these same types of assets, which raises the question of whether they now share a common AI exposure with the remainder of the portfolio. Because private assets are not valued at a frequency that permits direct estimation, we estimate exposures for publicly listed real-asset comparables, relative to a market benchmark from which the growing AI complex has been removed, and then map the estimates to private holdings by industry. We find that AI-related equity returns are not well described by a single factor, and we distinguish AI demand (the adoption of AI products) from the capital expenditure related to the AI buildout. The public comparables in power, midstream energy, electrical equipment, construction, and data centers have recently acquired a sizable exposure to the buildout, with no corresponding exposure to demand. For the private book, we separate a *risk-shift*, the higher implied sensitivity of assets already held, from a *mix-shift*, the reallocation of the portfolio toward the exposed sectors. The risk-shift has occurred and is inherited through existing holdings without any reallocation. The mix-shift is modest to date but is likely to grow and to compound the risk-shift, given that uncalled capital is substantial for newer funds that tilt toward the exposed sectors. The resulting exposure is best evaluated at the level of the total portfolio, alongside the related exposures held in public equity and credit.

Keywords: artificial intelligence; real assets; infrastructure; private markets; factor exposure; data centers; power

*All authors are with the Kenan-Flagler Business School at the University of North Carolina at Chapel Hill and the Institute for Private Capital (IPC). We thank IPC's Real Assets Research Initiative advisory board for insights and feedback (Albert Lee, Andrew Hoffman, Monica Huffer, Bill Kieser, Nick Russell, Shawn Gilligan, Sam Spencer, Louis Yu, and Willem Vlaming) as well as StepStone for providing data on real-asset fund holdings.

1 Introduction

Institutional investors traditionally allocate to private real assets, such as infrastructure, energy and power, and real estate, for reasons that are well understood. The assets generate long-duration cash flows that are frequently contracted or regulated, those cash flows are often linked to inflation, and their measured sensitivity to public equity markets has historically been low.[1] Global allocations to private real assets have increased from less than $200 billion in 2005 to $2.7 trillion in 2026.[2] Real assets are held, at least in part, as diversifiers, and portfolio construction has generally treated them as such.

The capital expenditure supporting artificial intelligence complicates this view, given the considerable scale of the program. The capital expenditure of the largest technology companies exceeded $400 billion in 2025 and is expected to increase by a further 75 percent in 2026. By way of context, the capital expenditure of just five firms now exceeds global investment in oil and natural gas production (International Energy Agency, 2026).[3] Projections of cumulative data center investment through 2030 range from $5 trillion to $7 trillion (Noffsinger et al., 2025; Bank of England, 2025); Van Nieuwerburgh (2026) estimates that the U.S. data center capacity likely to be completed by 2032 implies investment of roughly $9 trillion over 2025–2032, or 3.2 percent of GDP per year, a share that exceeds the canal, railroad, electrification, highway, and telecommunications booms. AI-related investment already accounts for a larger share of U.S. GDP growth than did information technology investment at the peak of the earlier technology cycle (Rubinton and Patro, 2026). While the majority of this expenditure is for computing equipment, roughly 40 percent is directed toward physical capacity, namely land and site development, power generation and transmission, and cooling and electrical equipment (Noffsinger et al., 2025), and data centers are expected to account for nearly half of the growth in U.S. electricity demand through 2030 (International Energy Agency, 2025).

[1] See, for example, Brown et al. (2025) and Brown et al. (2026).

[2] Source: MSCI Private Capital Universe.

[3] Company guidance is consistent with these figures. Based on 10-Q filings and earnings calls as of the second-quarter 2026 reporting season, the combined 2026 capital expenditure guidance of Alphabet, Amazon, Meta, and Microsoft is approximately $748 billion (Bloomberg). Van Nieuwerburgh (2026) reports roughly $800 billion for the same four firms together with Oracle, and notes that 2026 is the first year in which the capital expenditure of the five is projected to exceed their combined operating cash flow. For the eleven firms of the demand complex with analyst coverage, I/B/E/S fiscal-2026 consensus capital expenditures total approximately $630 billion (Appendix Table A12), with the difference reflecting fiscal-year alignment and the distinction between company guidance and analyst consensus forecasts.

These magnitudes are linked directly to real assets. Power generation, transmission, pipelines, and data centers are the types of assets that private real-asset portfolios typically hold, and private capital is a principal source of financing for the buildout. Of the nearly $3 trillion of data center expenditure expected through 2028, roughly half must be financed externally, and private capital is expected to supply the majority of that amount (Morgan Stanley Research, 2025). As the capital expenditure of the hyperscalers outgrows their internal cash flow, the marginal financing of the buildout is shifting toward leases, joint ventures, project debt, private credit, and securitization. These structures separate the users of computing capacity from the owners of the underlying assets and place the ownership of data centers and power with the developers, infrastructure funds, and REITs that real-asset investors hold (Van Nieuwerburgh, 2026).[4] To the extent that the value of these assets now depends on the continuation of the capital expenditure program, real assets share a common source of risk with the technology holdings that represent a growing share of public equity portfolios and, increasingly, with the credit instruments that finance the buildout (Aldasoro et al., 2026). The goal of this paper is to examine whether private real assets (and especially infrastructure) continue to serve as relatively diversified, long-duration inflation hedges in the presence of a common AI exposure. We focus primarily on diversification by asking how much of a private real-asset book now depends on the AI buildout and whether that dependence arises through a change in the risk of assets already held, through a reallocation of capital, or both.

The question is particularly pressing given the potential change in overall portfolio risks. For example, consider investors making allocations based on incremental risk contributions to a diversified portfolio, perhaps using a total portfolio approach (Ang, 2014; CAIA Association, 2024).[5] In this context, investors estimate how much exposure the fund as a whole has to a given source of risk, regardless of the asset class in which that exposure resides. An investor taking this view of the AI capital expenditure program requires the exposure of public equity, of credit, and of real assets, measured on a consistent basis. Whereas the exposure of public equity can be esti-

[4]For example, in October 2025, a consortium that includes BlackRock's Global Infrastructure Partners agreed to acquire Aligned Data Centers (Aligned Data Centers, 2025), and Meta placed its Hyperion campus in a joint venture that is majority-owned by funds managed by Blue Owl Capital (Meta Platforms, 2025); the Hyperion vehicle financed roughly $27 billion of a $30 billion asset with investment-grade project debt, a debt-to-asset ratio near 90 percent, against lease and residual-value commitments from Meta that do not appear on Meta's balance sheet until they take effect (Van Nieuwerburgh, 2026).

[5]CalPERS, the largest public pension fund in the United States, adopted a total portfolio approach in November 2025, effective July 2026 (CalPERS, 2025).

mated directly, that of private real assets cannot, and an implied measure of the latter is a primary contribution of this paper.

A natural starting point would be a single AI factor against which real assets could be evaluated. However, we show that AI-related equity returns do not share a single common component of that kind. We instead distinguish three layers, each with a different source of cash flows. The *demand* complex comprises the firms that sell AI products and services, namely the hyperscale platforms and the providers of AI software, and the valuations of these firms depend on adoption and monetization. The *buildout* tier comprises the firms that supply the equipment with which computing capacity is expanded, namely semiconductors, networking, servers, and the electrical and thermal equipment inside the data center; the valuations of these firms depend on the pace of capital expenditure. We construct a return factor from each of these two layers.

The third layer comprises real-asset firms and is our primary focus. These are the firms that own, operate, or construct the physical capacity on which the first two layers depend. These include independent power producers and utilities that generate electricity, midstream energy operators that transport natural gas used in generation, REITs that own data centers, manufacturers of heavy electrical equipment, and construction contractors building sites. The link between these firms and AI is indirect, in that they sell neither AI services nor computing equipment. Instead, they sell power, space, heavy equipment, and construction services to the firms that do (as well as to many other parts of the industrial economy). Two channels connect their value to AI. The first is the buildout channel, as their revenues and growth prospects depend on the pace at which capacity is added (irrespective of whether the AI services sold from that capacity are profitable). The second is the demand channel, as the value of capacity will eventually depend on the services that it supports and thus on the commercial success of the technology (e.g., through the renewal of a lease or a power contract and the credit quality of the counterparty).

Several of these sectors are also those in which private real-asset investors have traditionally invested. Power generation, utilities, midstream energy, and communications infrastructure are core holdings of private infrastructure funds, data centers have become an established property type in private real estate and infrastructure programs, and these sectors account for a sizable portion of the private holdings that we observe. The listed firms in these sectors serve as our

public comparables.[6] We deliberately exclude the real-asset firms from both factors, so that any exposure we find is not present by construction.

The distinction between demand and buildout is substantive, as an efficiency gain that allows the platforms to deliver more with less capacity is favorable for demand and adverse for the buildout. The two factors can therefore respond in opposite directions to the same information. For example, consider the release of the DeepSeek model in January 2025, which was interpreted as evidence that AI services could be delivered with considerably less computing capacity. On January 27, 2025, NVIDIA declined by 17 percent and the independent power producer Vistra by 28 percent, whereas the responses of the hyperscale platforms were mixed and small. More generally, the weekly returns of the two AI factors, demand and buildout, share roughly half of their variation, and roughly one-third once their common exposure to the ex-AI market is removed. As real assets represent physical capacity, the buildout factor is the natural candidate for their exposure, and the absence of a single AI factor is another of our primary findings.

The growth of the AI complex also bears on measurement. The combined market capitalization of the firms that we classify as AI-related increases from $8.3 trillion at the end of the first quarter of 2022 to $19.1 trillion four years later, or from 15.5 percent to 26.1 percent of total U.S. market capitalization. The value-weighted average market beta is equal to one by construction, such that, as a high-beta AI complex becomes a larger share of the index, the measured betas of the remainder of the market decline even where the underlying economic exposures are unchanged. We therefore measure all exposures relative to an ex-AI market benchmark.

Two distinctions organize the paper, and we maintain them throughout. The first is between public market comparables and private real-asset investments. Private assets are appraised infrequently, and their reported returns do not permit reliable estimation of factor exposures over the short and rapidly changing horizon that we study. All of our regression estimates are therefore obtained from publicly listed firms, grouped into real-asset sector portfolios, and we refer to these as the public comparables. Statements about private real assets are then obtained by applying the public market estimates to the observed point-in-time sector composition of private holdings,

[6] A private power plant or data center is appraised infrequently, whereas a listed independent power producer or data center REIT, which holds assets of the same type, is priced continuously. In some cases, the assets are the same, as QTS and CyrusOne, two of the largest private data center platforms, were listed REITs until their acquisition by private capital in 2021 and 2022 (Section 3).

which we obtain from investment-level data on private funds provided by StepStone, and we refer to the result as the *implied* exposure of the private assets. The mapping assumes that a private asset in a given sector experiences the same change in sensitivity as its listed comparables, and it abstracts from differences in leverage between listed firms and private funds.

The second distinction concerns the source of any change in the implied exposure of the private book. A *risk-shift* is a change in the AI sensitivity of the sectors that the book holds, with the composition of the book held fixed. The risk-shift is identified from the public comparables. A *mix-shift* is a change in the composition of private holdings toward, or away from, the sectors that carry AI exposure, with sensitivities held fixed. The mix-shift is identified from the private holdings data.

We document three main findings. The first two concern the public comparables, and the third concerns the private book. First, in public markets, the listed comparables for the assets held in private real-asset portfolios acquire a sizable exposure to the AI buildout. Specifically, buildout exposures increase significantly across public real-asset sectors, and the increases are economically large for independent power producers, electric utilities, heavy electrical equipment, and data center REITs, among others. The timing of the shift is also informative, as the corresponding estimates before 2025 are close to zero, which suggests that the repricing coincides with the period in which the scale of the capital expenditure program became evident rather than with the arrival of the technology. Further, none of the public real-asset portfolios acquires a significant positive exposure to AI demand. The evidence points to the physical buildout of capacity, rather than the adoption of AI products, as the relevant source of risk for real assets.

Second, and again in public markets, the new exposure partly displaces the old. Across the exposed sectors, the coefficient on the ex-AI market portfolio declines, on average, by a little more than half of the increase in the buildout coefficient. However, the displacement is not uniform. For heavy electrical equipment and for construction and engineering, the decline in the ex-AI coefficient is roughly as large as the increase in the buildout coefficient; for these sectors, a conventional market beta would register little change, because the index against which it is measured has moved in the same direction. For utilities and independent power producers, in contrast, the decline is small, and the buildout exposure is better viewed as an addition to systematic risk. We interpret both statements with caution and subsequently discuss the relevant caveats.

Third, turning to private real assets, the implied buildout beta of the stock of private holdings increases in 2025. Of this increase, the larger portion is derived from an implied risk-shift, inferred from the public comparables, and the smaller portion from an observed mix-shift in private holdings. The evidence for the risk-shift is strong: it derives from public market estimates that are statistically significant, that survive adjustment for multiple testing in five of the eight sectors, and that show no evidence of pre-existing trends. It applies to positions that private funds already hold, so no reallocation (mix-shift) is required for a private fund to have a substantial increase in AI buildout risk. The mix-shift, in contrast, is at an early stage, and we expect its contribution to grow for several reasons. First, the private holdings data we observe are reported with a lag, such that investment activity in the most recent years is mechanically understated. Second, capital committed to private funds is drawn over a period of several years, which implies that recent commitments have, for the most part, yet to appear in holdings. Third, the direction of new private investment is evident in the data. For example, data center transactions increase from near zero before 2017 to roughly 50 in 2025. Furthermore, the capital that new private funds have yet to deploy, together with the capital now being raised, is likely tilted toward the exposed sectors by a wider margin than the capital already invested. Importantly, the two components compound. New private capital is entering the sectors in which the public comparables exhibit the largest increases in buildout sensitivity, such that a growing mix-shift would be applied to coefficients that the risk-shift has already raised.

Relation to the literature

The paper draws on four strands of work. The first concerns benchmarks. We rely on the observation in Roll (1977) that beta is a property of the pairing between a security and an index rather than of the security alone, such that it moves when the composition of the index moves; see also Stambaugh (1982), Kandel and Stambaugh (1987), Shanken (1987), and Roll and Ross (1994) on the sensitivity of inference to the proxy, and Blume (1971) and Fama and French (1997) on the stability of estimated betas. The second concerns the measurement of risk in private real assets. Appraisal-based returns are smoothed, which understates volatility and comovement with public markets (Geltner, 1991, 1993; Getmansky et al., 2004; Couts et al., 2024; Hu et al., 2025), and the available corrections are not designed to detect a change in exposure over a short window. We therefore

estimate exposures on listed firms that hold the same types of assets, consistent with the evidence that listed and private real estate share a common value beyond the appraisal lag (Pagliari et al., 2005; Horrigan et al., 2009; Hoesli and Oikarinen, 2012; Ling and Naranjo, 2015), and we map the estimates to private holdings following Dittmar and Lundblad (2017).[7] The third concerns the role of real assets as diversifiers and inflation hedges (Fama and Schwert, 1977; Hartzell et al., 1987; Bekaert and Wang, 2010; Newell et al., 2011; Rödel and Rothballer, 2012; Bird et al., 2014; Pagliari, 2017; Andonov et al., 2021); we show that the systematic risk of the underlying assets is subject to change, independent of the vehicle through which those assets are held. The fourth concerns the economics of investment booms. The risk associated with a new technology becomes systematic as adoption on a large scale becomes likely (Pástor and Veronesi, 2009), the systematic risk of a firm changes as it acquires growth options (Berk et al., 1999; Carlson et al., 2004), and the owners of long-lived capacity that takes time to build are exposed to the supply response (Kydland and Prescott, 1982; Greenwood and Hanson, 2015). Existing work on AI concerns the firms that employ the technology (Eisfeldt et al., 2026; Babina et al., 2024; Acemoglu, 2025) and, more recently, the financing, macroeconomic, and local consequences of the buildout (Aldasoro et al., 2026; Rungcharoenkitkul, 2026; Caballero, 2026; Wachter and Wachter, 2026; Kay et al., 2026; Alvarez et al., 2026). Closest to our work is Van Nieuwerburgh (2026), who quantifies the scale of the buildout, describes the capital stack through which data centers and power are increasingly owned and financed outside the hyperscalers' balance sheets, and documents that the market beta of listed data center REITs has risen from roughly one-half to roughly one since 2018. We take the ownership structure he describes as the reason that the exposure resides in real-asset portfolios, and we extend the asset pricing evidence in two directions: we separate the buildout from AI demand, which attributes the repricing of data center REITs and of the other exposed sectors to the pace of capital expenditure specifically, and we carry the public estimates to the composition of private holdings. To our knowledge, no research measures the exposure of private real assets to the buildout.

[7] The corresponding evidence for infrastructure is thinner. Bianchi et al. (2014) and Blanc-Brude et al. (2017) find that broad listed infrastructure indices are largely spanned by conventional asset classes and equity factors, which is why we construct narrow sector portfolios; see also Bustos and DeMond (2025).

2 The Anatomy of AI Exposure

The classification on which the remainder of the paper rests sorts AI-related firms, and the real-asset firms that may be exposed to them, into three layers according to the source of their cash flows. The sorting is based on what a firm sells, and to whom, rather than on the behavior of its stock price.

Layer one: the demand complex: The first layer, the *demand* complex, comprises the firms that sell AI products and services: the hyperscale platforms and the providers of AI software and applications. The value of these firms depends on whether AI products are adopted and, in turn, generate revenue. The layer represents $8.4 trillion of market capitalization in early 2026, and its value-weighted return is the *AI demand* factor. The layer is concentrated, as its five largest constituents account for 96 percent of its market capitalization in early 2026, such that the value-weighted demand factor is, to a close approximation, a portfolio of the hyperscale platforms.

Layer two: the buildout tier: The second layer, the *buildout* tier, comprises the firms that supply the "picks and shovels": semiconductors, networking equipment, servers, data center hardware, and electrical and thermal equipment. The value of these firms depends on the continuation of capital expenditure, which is an economic exposure distinct from whether AI products are taken up. An efficiency gain that allows the platforms to deliver the same services with less capacity is favorable for the demand complex and adverse for the buildout tier. The tier represents $8.7 trillion of market capitalization in early 2026, and its value-weighted return is the *AI buildout* factor, the primary object of interest in this paper. The buildout tier is also concentrated, as NVIDIA and Broadcom represent 48 and 17 percent of its market capitalization in early 2026, respectively.[8]

Layer three: the public real-asset comparables: The final layer comprises the public real-asset comparables, the listed firms that own, operate, construct, or supply the physical assets in which private real-asset funds invest. These firms are the closest public analog to a private infrastructure portfolio. They are the test assets of the paper, and they are deliberately excluded from both fac-

[8] To explore robustness, we consider alternative specifications that employ equal-weighted and ex-NVIDIA buildout factors; see Section 5.

tors. We form fifteen sector portfolios and group them by the channel through which the AI buildout would be expected to reach them. The power and fuel channel comprises independent power producers and renewables, electric utilities, regulated utilities other than electric, midstream energy, and integrated oil and gas. The equipment and shell channel comprises data center REITs, heavy electrical equipment, and electrical components and equipment. The construction and materials channel comprises construction and engineering, building products, steel, copper, and aluminum. Tower and fiber REITs represent the connectivity channel, and traditional REITs are included as a real-asset sector for which we expect no direct link to the buildout; Section 3 describes the construction of the portfolios.

The one-home rule and the ex-AI market

Each firm is assigned to one of three layers or to a separate mixed/other AI category. Firms on the demand side enter the demand factor, equipment suppliers enter the buildout factor, and asset-heavy firms that own or operate physical capacity are treated as comparables. A small number of firms whose cash flows span more than one layer (Tesla, Oracle, CoreWeave, Nebius, and Eaton, representing roughly $2 trillion of market capitalization in early 2026) are assigned to neither factor; they are removed from the market benchmark together with the constituents of the two factors, but remain eligible for the broader "All AI Firms" group in the capital expenditure analysis when sufficient forecast data are available. If data center REITs were constituents of the buildout factor, a finding that data center REITs load on that factor would be circular. The market benchmark used throughout the paper is accordingly the *ex-AI* market, defined as the value-weighted market excluding the demand complex, the buildout tier, and the mixed names, and Table 1 lists the constituents of each group.

Constituent classifications are fixed over the sample, but index membership reflects the availability of each firm in the public market. Firms therefore enter an index only once they are publicly traded and observable in CRSP; later listings and spinoffs are not backfilled into earlier periods. For example, Arm enters in September 2023, Astera Labs in March 2024, GE Vernova in April 2024, Nebius in October 2024, and CoreWeave in March 2025. Consequently, the composition of the factors changes as new constituents become publicly traded, and the pre-2019 buildout factor is predominantly a semiconductor portfolio. We do not impose any mechanical floor on market

capitalization.

TABLE 1: AI Factor Constituents

AI Category	Companies
Buildout	NVIDIA, AMD, Broadcom, TSMC, Marvell, Micron, Credo, Ambarella, CEVA, Arm, Astera Labs, ASML, Applied Materials, Lam Research, KLA, Cadence, Synopsys, Super Micro Computer, Dell, HPE, Arista Networks, Ciena, Coherent, Lumentum, Applied Optoelectronics, Vertiv, GE Vernova, Celestica, Flex, Jabil, Modine
Demand	Microsoft, Amazon, Alphabet, Meta, Palantir, C3.ai, SoundHound AI, BigBear.ai, Veritone, UiPath, ServiceNow, Snowflake, AppLovin, Recursion Pharmaceuticals, Schrödinger, Absci, Certara, Twist Bioscience, Tempus AI, Mobileye, Symbotic, Aurora
Mixed/Other	Tesla, Oracle, CoreWeave, Nebius, Eaton

The two factors are distinct in the data as well. Table 2 reports that, since 2023, their weekly returns share roughly half of their variation, and that the shared portion falls to roughly a third once each factor's exposure to the ex-AI market is removed. Nearly two-thirds of their non-market variation is therefore distinct, which is the basis for treating the buildout and demand as separate exposures throughout the paper.

TABLE 2: Shared Variation Between AI Buildout and Demand Factors

Period	N	Shared Variation	Shared Variation, Ex-AI Adjusted
2023	52	42.5%	33.1%
2024	52	49.8%	39.8%
2025–2026Q1	65	56.1%	38.0%
2023–2026Q1	169	49.5%	36.3%

Shared variation is the squared correlation between weekly excess returns on the AI buildout and AI demand factors. Ex-AI adjusted shared variation is the squared correlation between residual returns after separately removing each factor's exposure to the ex-AI market.

3 Data and Measurement

The paper relies on two sources of data. The first comprises the public equity returns from which we build the AI factors and the real-asset comparables. The second comprises the private holdings to which the public market estimates are later applied.

3.1 Public equity data

The analysis in this paper involves several bespoke indices using weekly CRSP returns data combined with Global Industry Classification Standard (GICS) data from Compustat to create specialized exposure portfolios. We retain common equities issued by corporations and REITs with primary listings on the NYSE, Nasdaq, or NYSE American, requiring nonmissing returns and positive lagged market capitalization. The sample consists of weekly Friday-to-Friday returns running from the first week of 2010 to the end of the first quarter of 2026.

3.2 Real-asset portfolio construction

The publicly listed real-asset comparables are organized into the fifteen sector portfolios. Each is a value-weighted portfolio of the listed firms in the corresponding GICS sub-industries, and Appendix A lists the sub-industries assigned to each portfolio. We build the portfolios from their constituents rather than rely on published sector indices for two reasons. First, construction at the level of the constituent allows us to enforce the one-home rule under which no firm appears both in a factor and in a comparable portfolio. Second, published sector series are not consistently defined over our sample. In particular, the GICS structure was revised in March 2023, when Data Center REITs became a distinct sub-industry, having previously been classified within Specialized REITs together with tower operators and other property types. We apply each firm's latest available Compustat GICS classification retrospectively, supplemented by explicit assignments for the data center and tower/fiber REIT portfolios, as detailed in Appendix A.

The data center REIT portfolio warrants a further comment. The number of listed pure-play data center REITs in the United States declined from six in 2016 to two by the end of 2022. Our broader comparable portfolio includes Iron Mountain, bringing the number of constituents to three at year-end 2022.[9] Fermi, a development-stage company organized as a REIT, enters the portfolio upon its listing in October 2025, which brings the number of constituents to four at year-end 2025. The thinning of the listed cohort has two implications. It is a measurement concern, as the portfolio rests on a small number of large firms in the latter part of the sample, and the

[9] DuPont Fabros was acquired by Digital Realty in 2017 and CoreSite by American Tower in 2021, in both cases by a listed acquirer, whereas QTS in 2021 and CyrusOne in 2022 were acquired by private capital; Switch, a listed data center operator that was not organized as a REIT, was also taken private in 2022.

remaining listed firms may not be fully representative of the single-tenant hyperscale facilities that private capital has financed. It is also consistent with the broader thesis of the paper, in that the marginal data center asset is increasingly built or acquired in private markets.

3.3 Private holdings: StepStone

Data on private real-asset investments are provided by StepStone. The data are recorded at the level of the individual investment, and we combine the private infrastructure and private real estate investment files into a single real-asset sample. Unless otherwise noted, the exposed shares that we report are measured against this combined sample, such that real estate other than data centers is included in the denominator and is assigned no buildout exposure. The extract was obtained on August 13, 2026.[10] For each investment, we observe the year of entry, the sector, the capital invested, the reported total value (the sum of realized proceeds and unrealized value), and, where available, an identifier for the fund through which the investment was made. The fund-level sample contains 12,178 observed investments with entry dates from 1991 through 2025; among observations with an available fund identifier, 1,058 distinct funds are represented. Because the fund-level analysis retains the full observed history of funds active in 2010 or later, it includes some investments entered before 2010. The deal-level analysis instead restricts individual investment entry dates to 2010–2025 and contains 11,139 investments.

Each private investment is assigned to one of the public sector portfolios on the basis of its 8-digit GICS sub-industry sector classification, which provides the link between the private holdings and the public market estimates. Private investments in sectors for which we have no public comparable are retained in the denominator of all portfolio weights and are assigned no change in buildout exposure.

Two features of the data are important for what follows. First, investments enter the data with a reporting lag such that private investment activity in the most recent years is understated. We return to the implications of the lag for the measurement of the mix-shift later. Second, the data do not include a time series of interim valuations. We therefore construct a synthetic net asset value (NAV) path for each investment over its observed life using invested capital at entry

[10] We restrict the analysis to closed-end funds and exclude evergreen (quasi-open-end) vehicles and their investments throughout; because the evergreen vehicles that we observe hold a larger share of their value in the exposed sectors, the exclusion errs toward understating the exposure of the private book.

and terminal value information observed at exit for realized investments or at the most recent report date for active investments. The synthetic values serve only as weights for private-market exposure. They are not used to estimate returns or risk exposures, all of which are obtained from the public comparables. A detailed description of the synthetic NAV process can be found in Appendix F.

The fund-level analysis in Sections 6.4 and 6.5 requires complete fund holdings. We identify 159 funds for which the reported number of investments is known and all investments are observed, and refer to these as the clean sample. For an additional set of partially observed funds, we reconstruct missing fund memberships using a multinomial logit model estimated on the clean sample. Reconstructed funds are retained only if they satisfy a minimum fit criterion calibrated on observed clean funds. Because the set of retained reconstructed funds varies across reconstructions, the analysis is conducted separately on the top 20 reconstructed panels, each combining the clean funds with the reconstructed funds retained in that panel, and results are averaged across these panels. Appendix E describes the reconstruction procedure and reports its out-of-sample accuracy. The clean funds are smaller, more heavily weighted toward real estate, and less exposed to the sectors of interest than the funds that require reconstruction, such that estimates on the clean sample alone are best read as a lower bound. The 2025 book-level buildout exposure change is 0.118 on the clean sample and 0.186 across the reconstructed panels, and we report both samples where the fund level is the unit of analysis.

4 The Moving Yardstick

With the factors and the comparable portfolios in hand, the natural next step would be to estimate the sensitivity of each real-asset sector to the market and to the AI factors. Before doing so, we must address a measurement problem that affects any such estimate. Every beta in this paper is measured against a market portfolio, and since 2022 the composition of that portfolio has changed faster than the economics of most of the firms in it. A single-benchmark beta cannot serve as a measure of changing risk under those conditions. Below we explain why, what the arithmetic implies for the patterns in Figure 1, and how the ex-AI benchmark and our multi-factor regressions address the problem. The argument is a corollary of an observation in Roll (1977).

The practical content can be stated without the algebra. A market beta compares an asset

with an index, and the index has changed: firms with a common AI exposure have grown from roughly one-sixth to roughly one-quarter of its value. Because the betas of all firms in an index must average to one, the betas of the firms outside the AI complex decline as the complex grows, even when nothing about those firms has changed. For the same reason, a real-asset sector with a stable market beta can also have changed, as it may have acquired AI exposure at the same pace as the index. To examine this empirically, we measure market exposure relative to an index that excludes the AI complex, and we measure AI exposure separately and directly.

Roll's critique is remembered for one conclusion: the true market portfolio is unobservable, so any test of the two-parameter model is a joint test of the theory and of the efficiency of the proxy, and thus the theory is untestable in practice. That conclusion concerns the pricing of assets. We are not testing an asset pricing model; we use betas as descriptive statistics of comovement, and for that purpose the efficiency of the proxy is irrelevant.[11] A second observation in the same paper is the one that matters here. Roll noted that beta is a relative measure, so that two investors holding different "market" portfolios, both efficient, might assign the same security betas of 1.5 and 0.5. He then made the less obvious remark that if both investors increased the proportion the security represents in their respective markets, its beta would change, and could rise for one investor and fall for the other. He added that most reasonable proxies are so highly correlated that composition appears unimportant, "whereas it can cause quite different inferences," and that the profession's early reading of the cross-sectional evidence would have been different had the index's composition "been correctly perceived as the critical variable."

Beta is a property of the pairing between an asset and an index, not of the asset alone. When the composition of the index changes, every beta measured against it changes, whether or not anything about the underlying assets has changed. The U.S. equity market since 2022 is a direct instance: on the figures in Section 1, a set of firms with a common exposure grew from roughly one-sixth of the index to roughly one-quarter of it in four years, and every beta measured against the index moved for that reason alone; the remainder of this section makes the statement algebraically exact.

[11] The pricing question is taken up by Wachter and Wachter (2026), not in a factor model but in a two-sector equilibrium in which the hyperscalers' capital expenditure reveals the size of a rare productivity boom, and in which both the risk-free rate and the equity premium rise as the AI sector's share of the economy grows; our covariances describe the implications of the growing AI share from a different perspective.

Let $w_{i,t}$ denote market-capitalization weights and $r_{m,t} = \sum_i w_{i,t}\, r_{i,t}$ the market return. For betas estimated against that same market over any window,

$$\sum_i w_i\, \beta_i \;=\; \frac{\mathrm{Cov}(\sum_i w_i r_i,\; r_m)}{\mathrm{Var}(r_m)} \;=\; \frac{\mathrm{Var}(r_m)}{\mathrm{Var}(r_m)} \;=\; 1, \tag{1}$$

exactly, up to the drift of weights within the window. The value-weighted average beta is one by construction, and three facts follow.

First, betas rank systematic risk; they do not measure its level. If systematic risk rises for the entire market at once, the cross-section of betas cannot register it. The increase instead appears in the variance of the market.

Second, beta drift is zero-sum. Partition the market into an AI complex with weight w_A and average beta $\bar{\beta}_A$ and the remainder, so that

$$1 \;=\; w_A\, \bar{\beta}_A + (1 - w_A)\, \bar{\beta}_{\text{rest}} \qquad \Longrightarrow \qquad \bar{\beta}_{\text{rest}} \;=\; \frac{1 - w_A\, \bar{\beta}_A}{1 - w_A}. \tag{2}$$

With the AI complex at 15.5 percent of the market and an average beta of 1.3, the remainder averages 0.94; at 26.1 percent and the same beta, the complement is forced down to 0.89. A decline in the beta of relatively non-AI sectors such as health care and consumer staples is an artifact of this renormalization—an accounting consequence of the AI complex's growing weight rather than evidence that the defensive industries changed character.

Third, a flat beta path for real estate and infrastructure over a period in which the AI complex's share of the market rose by two-thirds is not a null result. Under the identity, standing still against a moving benchmark requires moving with it.

4.1 A two-factor decomposition of the measured beta

Suppose returns obey a two-factor structure,

$$r_{i,t} \;=\; a_i + \beta_i^O\, F_{O,t} + \beta_i^A\, F_{A,t} + \varepsilon_{i,t}, \tag{3}$$

in which F_O is an "old economy" factor and F_A an AI factor, orthogonalized so that $\mathrm{Cov}(F_O, F_A) = 0$, with variances σ_O^2 and σ_A^2. The market inherits the weighted loadings $b_O = \sum_i w_i \beta_i^O$ and $b_A =$

$\sum_i w_i \beta_i^A$; the growing AI content of the market—rising AI weights, rising AI valuations—is a rising b_A. Ignoring the idiosyncratic terms, the measured market beta of sector s is

$$\beta_s \;=\; \frac{\mathrm{Cov}(r_s, r_m)}{\mathrm{Var}(r_m)} \;=\; \frac{\beta_s^O \, b_O \, \sigma_O^2 \;+\; \beta_s^A \, b_A \, \sigma_A^2}{b_O^2 \, \sigma_O^2 \;+\; b_A^2 \, \sigma_A^2}. \tag{4}$$

Four results follow from (4), and together they account for the patterns in Figure 1.

Result 1 (the unexposed must fall). For a sector with no AI loading, $\beta_s^A = 0$, the numerator of (4) is fixed while the denominator rises in b_A. Its measured beta falls as the market's AI content rises.

Result 2 (flatness is a footprint). Hold a sector's structural loadings (β_s^O, β_s^A) fixed and differentiate (4) with respect to b_A:

$$\frac{\partial \beta_s}{\partial b_A} \;>\; 0 \qquad \Longleftrightarrow \qquad \beta_s^A \;>\; 2\, b_A \, \beta_s. \tag{5}$$

The threshold, which is derived in Appendix D, is demanding: holding the market's old-economy loading and factor variances fixed, a sector with a positive and constant AI loading can still see its measured market beta decline as b_A rises. More generally, a flat rolling market beta need not indicate stable underlying factor exposures, because changes in a sector's loadings can be offset by changes in the composition of the market benchmark.

Result 3 (the miniature-market case). If a sector's loadings are proportional to the market's ($\beta_s^O = c\, b_O$ and $\beta_s^A = c\, b_A$ for some constant c), then (4) collapses to $\beta_s = c$ in every period, regardless of composition. A sector whose AI exposure grows in proportion to the market's own displays *a constant beta while its risk is changing*.

Result 4 (the level of risk resides in the market variance). The quantity an allocator manages is the covariance with the equity book, $\beta_s \, \sigma_m^2$, or in volatility units $\beta_s \, \sigma_m$, rather than β_s alone. The same growth in AI content that renormalizes betas raises $\mathrm{Var}(r_m) = b_O^2 \sigma_O^2 + b_A^2 \sigma_A^2$ through b_A and concentrates it in the AI factor, whose share of market variance is $b_A^2 \sigma_A^2 / \mathrm{Var}(r_m)$. A flat beta

against a more volatile and more concentrated market is a rising risk contribution per dollar.

These results do not imply, however, that the current-market beta is the wrong number for pricing or risk-budgeting today; covariance with today's market is the basis of today's diversification argument. The claim is that the time series of single-benchmark betas cannot measure changing risk when the composition of the benchmark is itself the fastest-moving object in the system, such that the measurement of drift requires a benchmark that is held fixed.

4.2 Fixing the yardstick

Two remedies follow from this analysis. The first is an ex-AI market return: the capitalization-weighted market with the pre-specified AI complex removed and the remainder renormalized,

$$r_{m,t}^{\text{ex-AI}} = \frac{r_{m,t} - w_{A,t}\, r_{A,t}}{1 - w_{A,t}}, \tag{6}$$

where $w_{A,t}$ and $r_{A,t}$ are the weight and return of the complex: in our construction, the buildout tier, the demand complex, and the mixed names shown in Table 1 together. This benchmark is investable, is defined point-in-time, requires no base year, and membership is fixed by identity rather than by realized returns. Against it, the rolling beta of a real-asset comparable is free of the renormalization in (2). The gap between its path against the current market and its path against the ex-AI market is the compositional effect.

The second remedy is to name the factor: placing the AI returns on the right-hand side alongside the ex-AI market delivers the partial loading on each AI factor directly. Collinearity is modest because the ex-AI market excludes the complex, and no orthogonalization is required to interpret the coefficients. The buildout loading, and its change across windows, is the object of interest throughout the paper. We retain the current-market beta for contrast.

The construction also provides a placebo test. Health care and consumer staples should display stable ex-AI betas and AI loadings near zero. If the decomposition attributes their measured decline to the benchmark's drift, it is performing as intended.

4.3 Rolling betas against the market and the ex-AI market

To show the effect of this yardstick, we calculate 104-week rolling market betas for six of the real-asset portfolios, and for two placebo sectors, against the conventional market index and against the ex-AI market index. The results are shown in Figure 1, and the distinction is economically meaningful.[12] In each panel, the gap between the two lines is the portion of the change in measured beta that is attributable to the changing composition of the index rather than to the sector. The largest divergences between conventional and ex-AI betas appear in sectors with little AI exposure: the "placebo" sectors of health care and consumer staples and, among real assets, traditional REITs and tower and fiber REITs, with smaller gaps for regulated utilities and midstream energy. In each case, the conventional market beta falls relative to the ex-AI beta, consistent with a benchmark composition effect rather than the acquisition of meaningful AI exposure. By contrast, independent power producers and data center REITs exhibit barely perceptible differences between the two market beta measures, which is the pattern that Results 2 and 3 predict for sectors whose AI loading has risen with that of the market. The data center REIT panel is consistent with the evidence in Van Nieuwerburgh (2026), who reports that the 36-month rolling market beta of data center REITs has risen from roughly one-half to roughly one since 2018 and that their exposure to the non-AI component of the S&P 500 has risen alongside it; our rolling estimates likewise show the sector's ex-AI beta rising to above one in 2024, before declining over 2025 and early 2026. That rise is not in tension with the negative ex-AI exposure change that the synthetic-control estimates of Section 5 report for the same sector. The rolling beta is a level, estimated against the ex-AI market alone, whereas the synthetic-control estimate is a change relative to a counterfactual, with the buildout and demand factors on the right-hand side. Once the newly acquired buildout exposure and the change in the counterfactual are accounted for, the 2025 repricing of the sector loads on the buildout factor, and less of its systematic risk is attributed to the remainder of the market.

This motivates using the ex-AI market as the common market control in the subsequent analysis: changes in buildout exposure can then be interpreted as incremental changes in portfolio-specific risk rather than as a mechanical consequence of the growing weight of AI firms in the

[12] The gray area in each picture denotes the COVID period, and we encourage the reader to focus on the pre- and post-COVID differences separately from the unusual volatility of that moment.

aggregate market.

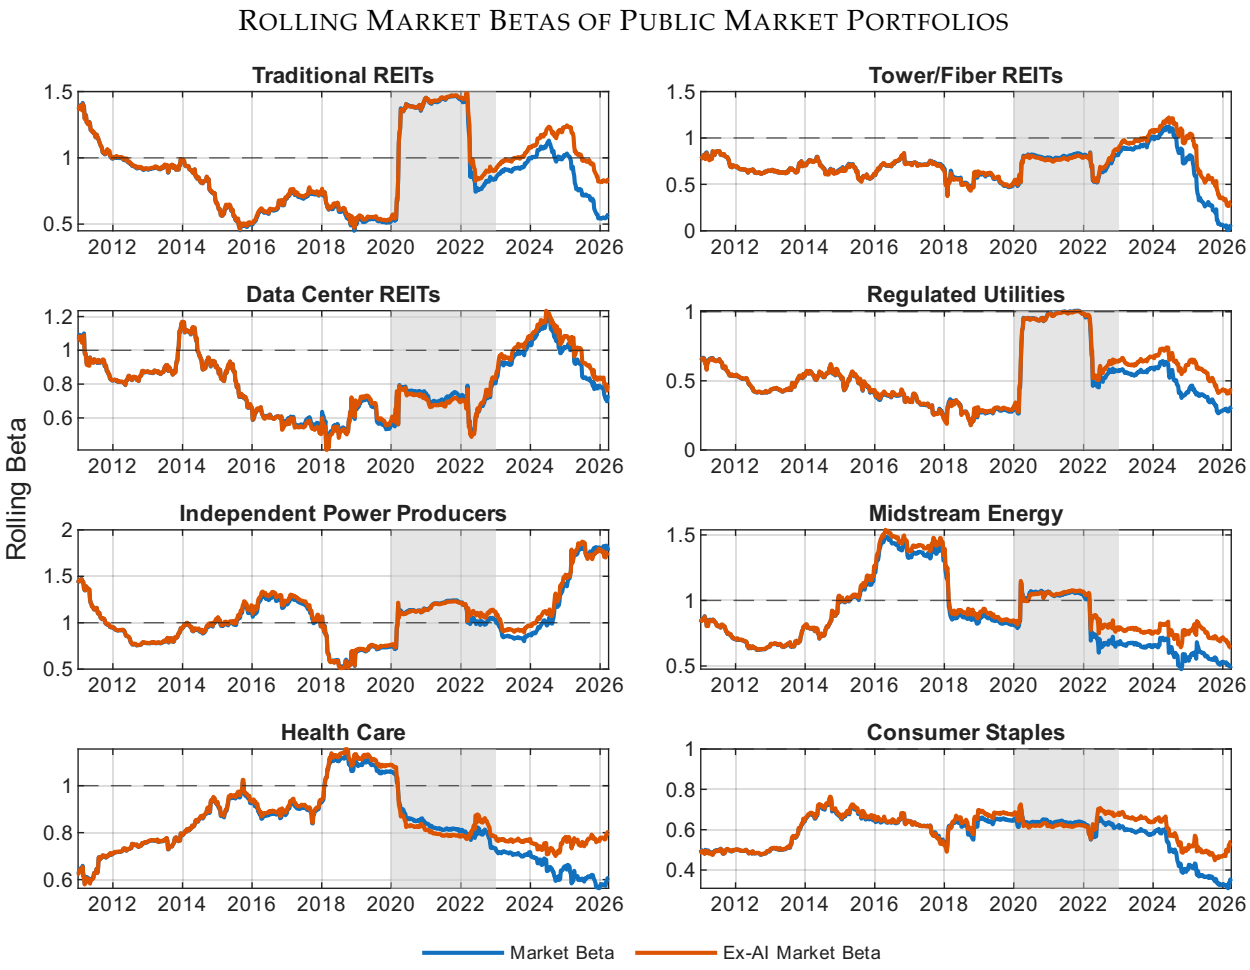


FIGURE 1: This figure shows 104-week rolling betas for each public-market portfolio, estimated using weekly returns against either the full market return or the ex-AI market return. Returns beginning in 2009 are used to initialize the rolling estimation windows, and the plotted sample runs from January 2011 through March 2026. The dashed horizontal line denotes a beta of one.

5 What Changed in the Public Comparables, and When

The preceding section establishes how exposure should be measured, and here we report what that measurement delivers for the public comparables. We proceed in three steps. We first state the regression that defines a sector's exposure to each factor. We then describe the comparison that separates a change attributable to the AI buildout from changes common to the market as a whole. Last, we report the results: which sectors acquire exposure, to which of the two AI factors, when, and whether the new exposure adds to or displaces exposure to the remainder of the market. All of these results concern publicly listed firms, and Section 6 later applies them to private holdings.

The four-factor regression we estimate is

$$r_{p,t} \;=\; \alpha_p \;+\; \beta_p^{M}\, r_{m,t}^{\text{ex-AI}} \;+\; \beta_p^{B}\, r_t^{B} \;+\; \beta_p^{D}\, r_t^{D} \;+\; \beta_p^{T}\, r_t^{T} \;+\; \varepsilon_{p,t}, \tag{7}$$

with $r_m^{\text{ex-AI}}$ the ex-AI market of (6), r^B buildout, r^D demand, and r^T the Treasury return.[13] In words, the weekly return on each real-asset sector portfolio is explained by the return on the market excluding the AI complex, the returns on the two AI factors, and the return on Treasuries. The buildout coefficient, β_p^B, measures the response of the real-asset sector, holding the remainder of the market fixed, to a one percent move in the firms that supply AI equipment in a given week, and the object of interest is the change in β_p^B across windows. Standard errors are heteroskedasticity- and autocorrelation-consistent (Newey and West, 1987) where a single regression is reported; inference on changes comes from the synthetic-control bootstrap described below.

Identification: synthetic-control difference-in-differences

Using the ex-AI market illustrates the moving benchmark problem, but a more informative look requires a counterfactual for how the risk exposures of each real-asset sector would have evolved absent the AI investment boom. We therefore employ a difference-in-differences (DiD) design to isolate changes in AI-related risk exposure from broader shifts in factor sensitivities across the market. Specifically, we follow a synthetic control approach as in Abadie et al. (2010). For each of the fifteen public real-asset portfolios, we construct a synthetic control as a weighted combination of plausible non-AI industries that closely matches the portfolio's pre-period sensitivities to four common factors: the ex-AI market, AI buildout, AI demand, and Treasury returns.[14] We estimate the synthetic control weights using factor betas from 2010–2017 and hold out 2018–2019 to check that the synthetic control still tracks the real portfolio during a later pre-AI period that was not used to build it; this period also serves as the common untreated baseline. We then compare the target-minus-synthetic difference in factor loadings after the AI boom with the corresponding difference in the baseline period. This difference-in-differences estimate measures the change in a

[13] As a control for movements in interest rates, all specifications include the weekly return on the Bloomberg U.S. Treasury 7–10 Year Total Return Index. The inclusion of this factor is not central to our results.

[14] The synthetic controls are built from a pool of eight broad and plausibly non-AI industry portfolios: consumer staples distribution and retail; food, beverage and tobacco; household and personal products; health care equipment and services; pharmaceuticals/biotech/life sciences; banks; financial services excluding mortgage REITs; and insurance. Any AI stocks that happen to lie within these portfolios are removed from the synthetic controls.

portfolio's factor exposure beyond that experienced by its synthetic counterfactual.

The synthetic control is designed to resemble the real portfolio as closely as possible before the AI boom, so that a later divergence is more plausibly interpreted as a change specific to that portfolio rather than a difference that was already present beforehand. The logic is that of any controlled comparison. A simple before-and-after comparison of a sector's buildout beta would attribute to AI any change in exposures that was common to the market over a period that included a pandemic and a sharp increase in interest rates. The synthetic control serves as the comparison group: a blend of sectors with no meaningful connection to AI, such as consumer staples, health care, banks, and insurers, weighted so as to behave like the real-asset sector before the AI period. We subtract the change in the exposure of the control from the change in the exposure of the sector, and what remains is the portion of the change that is specific to the sector.

Specifically, the difference-in-differences estimate for real-asset portfolio i and factor f, denoted by $\widehat{\Delta\beta}_{i,f}$, is given by

$$\widehat{\Delta\beta}_{i,f} = \left(\hat{\beta}^{\text{Target}}_{i,f,\,2025\text{–}2026Q1} - \hat{\beta}^{\text{Synthetic}}_{i,f,\,2025\text{–}2026Q1}\right) - \left(\hat{\beta}^{\text{Target}}_{i,f,\,2018\text{–}2019} - \hat{\beta}^{\text{Synthetic}}_{i,f,\,2018\text{–}2019}\right), \tag{8}$$

where f indexes the buildout, demand, ex-AI market, or Treasury factor.[15] For the remainder of the paper, we refer to this estimate as the sector's *exposure change* and, for the buildout factor, as its *buildout exposure change*. A buildout exposure change of 0.40 indicates that, relative to what would have been expected absent the AI buildout, the sector now moves by an additional 0.40 percent for each one percent move in the buildout factor.

Assuming that the effect of AI was zero prior to 2023, we estimate separate effects for 2023, 2024, and January 2025–March 2026, allowing us to trace when the divergence emerges rather than imposing a single post-treatment date. The effect is clearest and strongest in 2025, and we accordingly treat 2025 as the AI "event" throughout.

Buildout exposure

Table 3 reports the change in factor betas between the 2018–2019 baseline and January 2025–March

[15]Inference is based on a circular moving block bootstrap with 1,999 replications and a 13-week block length, with the synthetic control rebuilt in each draw to account for uncertainty in both returns and the construction of the synthetic counterfactual. The entire procedure is thus repeated on many resampled versions of the return history, and the dispersion of the resulting estimates provides the standard errors.

2026. The increase in buildout exposure is concentrated in sectors tied most directly to the physical AI investment cycle. Within the power and fuel complex, Independent Power Producers & Renewables exhibit a buildout exposure change of 0.66, Electric Utilities 0.38, Midstream Energy 0.26, and Regulated Utilities excluding Electric 0.24. The pattern also extends to equipment and physical infrastructure, with increases of 0.70 for Heavy Electrical Equipment and 0.48 for Data Center REITs, and to construction and materials, with increases of 0.50 for Construction & Engineering and 0.31 for Building Products. Five of these eight estimates are statistically greater than zero at the 95 percent confidence level or better based on unadjusted bootstrap inference. Heavy Electrical Equipment, Midstream Energy, and Building Products are significant at the 10 percent level, and five of the eight estimates remain significant at the 10 percent level after the Benjamini–Hochberg false-discovery-rate (FDR) adjustment across portfolios. For each of the eight portfolios, we also fail to reject parallel pretrends in buildout factor betas at the 90 percent confidence level using either annual or two-year pretrend tests. In plain terms, the significance tests ask whether increases of this size could have arisen by chance, and for five of the eight sectors the probability is below 5 percent. Because we examine fifteen portfolios, a few would appear significant by chance alone; the false-discovery adjustment guards against that possibility, and five of the eight survive it. The pretrend tests ask whether these sectors were already drifting away from their controls before the AI period, and we find no evidence that they were. As for magnitudes, a 10 percent decline in the buildout factor would now be associated with an additional decline of roughly 7 percent for listed independent power producers, 5 percent for data center REITs, and 4 percent for electric utilities, holding the ex-AI market fixed. These are large sensitivities for sectors that are traditionally held for their stability. The results indicate a broad repricing of buildout exposure across the power, infrastructure, and construction sectors most directly connected to the physical requirements of AI investment.

Demand exposure

The demand factor produces no systematic evidence of increased exposure, as none of the fifteen publicly listed real-asset portfolios exhibits a statistically significant positive demand exposure change. The contrast with the buildout results is sharp: the increase in AI-related risk among public real assets is associated specifically with the physical buildout of AI capacity rather than with

TABLE 3: Changes in the Factor Exposures of the Public Comparables

(Synthetic Control Difference-in-Differences Estimates)

Portfolio	**Buildout** $\Delta\beta$	**Demand** $\Delta\beta$	**Ex-AI Market** $\Delta\beta$	**Treasury** $\Delta\beta$
Heavy Electrical Equipment	0.698^{*}	0.483	-0.677	1.098
	(0.381)	(0.594)	(0.835)	(1.838)
IPP & Renewables	$0.664^{***\dagger}$	-0.103	-0.093	0.648
	(0.155)	(0.192)	(0.365)	(0.575)
Construction & Engineering	$0.496^{**\dagger}$	-0.294	-0.455	0.776
	(0.167)	(0.309)	(0.360)	(0.522)
Data Center REITs	$0.481^{**\dagger}$	-0.071	-0.311	-0.441
	(0.191)	(0.245)	(0.327)	(0.795)
Electric Utilities	$0.379^{***\dagger}$	-0.262	-0.045	0.382
	(0.089)	(0.181)	(0.285)	(0.411)
Building Products	0.305^{*}	-0.069	-0.245	0.053
	(0.145)	(0.313)	(0.372)	(0.477)
Midstream Energy	0.256^{*}	-0.368	-0.215	-0.152
	(0.122)	(0.202)	(0.435)	(0.567)
Regulated Utilities ex Electric	$0.244^{***\dagger}$	-0.247	0.038	0.152
	(0.086)	(0.167)	(0.287)	(0.337)
Integrated Oil & Gas	0.087	-0.502^{**}	-0.086	-0.602
	(0.132)	(0.191)	(0.422)	(0.755)
Electrical Components & Equipment	0.084	0.110	-0.123	0.560
	(0.121)	(0.161)	(0.232)	(0.483)
Traditional REITs	0.030	-0.127	0.389	-0.458
	(0.090)	(0.151)	(0.237)	(0.377)
Copper	-0.051	0.038	0.522	0.700
	(0.209)	(0.335)	(0.514)	(1.390)
Aluminum	0.005	0.119	0.081	0.969
	(0.226)	(0.325)	(0.442)	(1.091)
Tower/Fiber REITs	0.026	-0.111	-0.027	-0.011
	(0.188)	(0.297)	(0.504)	(0.901)
Steel	-0.169	-0.166	0.428	0.503
	(0.127)	(0.224)	(0.440)	(0.750)

This table reports the exposure change, $\Delta\beta$, for each public comparable portfolio and each factor, estimated as a synthetic control difference-in-differences in buildout, demand, ex-AI market, and Treasury factor betas. Synthetic controls are estimated using weekly data from 2010–2017, with 2018–2019 used as a held-out validation period and as the pre-period baseline. We omit 2020–2022 because the COVID-19 pandemic and the associated policy and market responses created unusually large disruptions that are difficult to disentangle from changes in underlying factor exposures. The exposure changes compare 2018–2019 with the post-period, January 2025–March 2026; observations from 2020–2024 are omitted from the pre/post comparison. Bootstrap standard errors are reported in parentheses. Inference is based on repeated resampling of blocks of weekly returns, with the synthetic control rebuilt in each draw to account for uncertainty in both returns and the construction of the synthetic counterfactual. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$ against a null of zero based on unadjusted bootstrap p-values. † denotes estimates that remain significant at the 10% level after Benjamini–Hochberg false-discovery-rate adjustment across portfolios. For all eight portfolios shown above the line, we fail to reject parallel pretrends in buildout factor betas at the 10% level using either annual or two-year pretrend tests. The full set of pretrend test results is reported in Appendix Table A2.

growth in demand for AI products and services. The practical implication is that the variable to monitor for a real-asset portfolio is the pace of capital expenditure commitments rather than the uptake of AI products.

Timing: 2023 versus 2025

For each of the fifteen sectors, buildout exposure changes for 2023, 2024, and 2025 are shown in Figure 2. Consistent with the results of Table 3, the majority of both large and statistically significant results are seen in the eight portfolios mentioned previously. However, a temporal disaggregation allows us to observe that the results are clearly strongest in 2025 and are often statistically indistinguishable from zero earlier. The timing is consistent with repricing as the scale of the capital expenditure program became evident, rather than at the point at which the technology became visible. Although we do not test that interpretation directly here, the capital expenditure forecasts of Section 6.6.2 provide the natural series against which to do so.

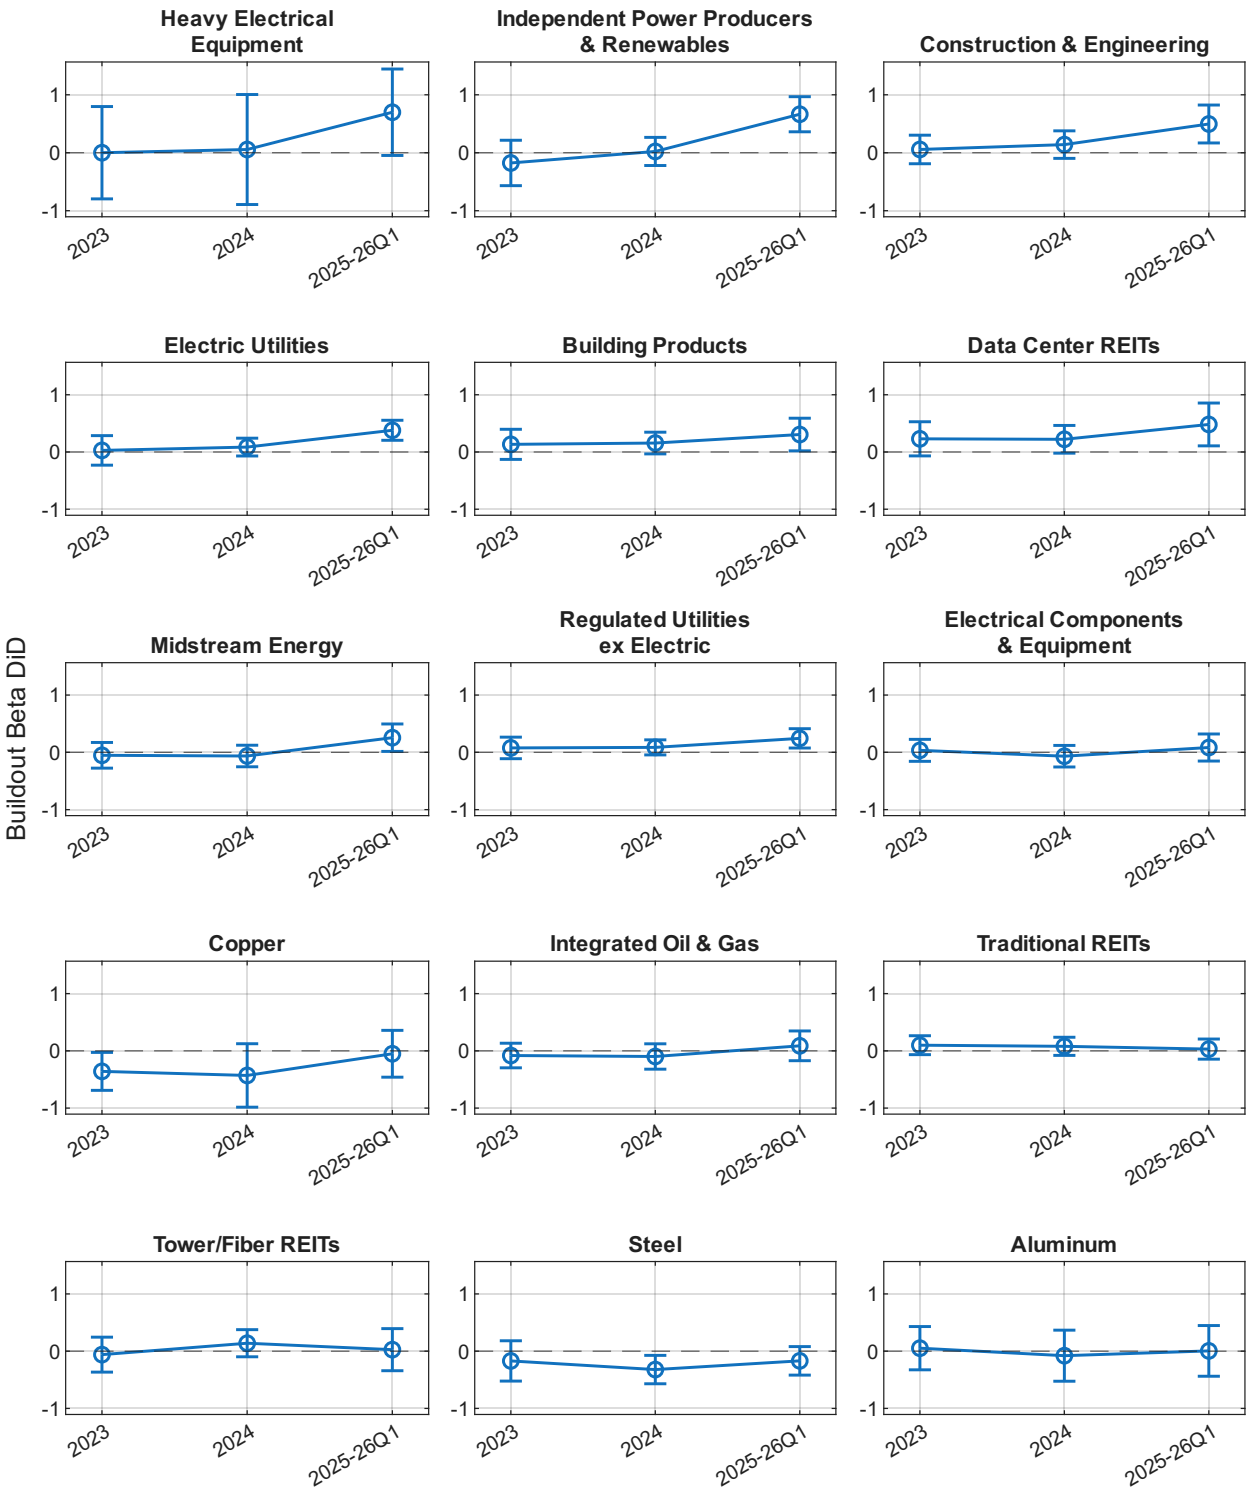


FIGURE 2: This figure shows synthetic control event study estimates of differential buildout-factor betas (buildout exposure changes) for each public market portfolio. Differential betas are measured relative to the portfolio's synthetic control and are reported for 2023, 2024, and January 2025–March 2026, relative to the 2018–2019 baseline. Vertical bars denote 95% confidence intervals.

The composition of systematic risk: migration rather than accumulation

The increase in buildout sensitivity in 2025 coincides with a decline in ex-AI market sensitivity across seven of the eight exposed sectors, consistent with a pronounced shift in the composition of systematic risk toward the AI buildout factor. Across the eight sectors, the average ex-AI shift is about -0.25 while the average buildout shift is about $+0.44$. This conforms to Result 2 from Section 4, as a single-benchmark beta would have obscured much of the change in the composition of systematic risk. The source of systematic risk in these sectors has therefore changed alongside its amount, as a portion of what was formerly general market exposure is now exposure to the AI buildout. A conventional market beta combines the two and therefore obscures this shift, which is the reason that a risk report based on a single market beta could have failed to flag the change. Two caveats apply to the displacement finding. The declines in the ex-AI coefficient are estimated with considerably less precision than the increases in the buildout coefficient, and none of the eight is individually significant at the 90 percent confidence level based on bootstrap inference. The displacement is therefore a statement about the average pattern across the exposed sectors rather than about any one of them, and the firmer conclusion is the one that the buildout coefficients support on their own: the exposed sectors now carry a source of systematic risk that they did not carry before 2025.

Robustness

The results of Table 3 are robust in direction and economic magnitude to alternative constructions of the counterfactual and to the weighting of the buildout factor, although the strength of the statistical evidence varies across specifications. We consider three alternative specifications, reported in Appendix C: excluding NVIDIA from the value-weighted AI buildout factor (Appendix Table A4); constructing an equal-weighted buildout factor (Appendix Table A5); and replacing the portfolio-specific synthetic controls with a common control comprising health care, financial, and consumer staples firms (Appendix Table A6).

Excluding NVIDIA from the value-weighted AI buildout factor leaves the main findings largely intact, with seven of the eight baseline-exposed sectors retaining FDR-adjusted significance and six satisfying both significance and pretrend criteria. Several estimated exposure changes also become larger, suggesting that the baseline findings are not driven by NVIDIA's dominant weight in the factor. Equal weighting likewise preserves the broad pattern of positive AI buildout ex-

posure changes but weakens statistical precision, leaving four sectors significant after multiple-testing adjustment, all of which also satisfy both pretrend tests. Replacing the portfolio-specific synthetic controls with a single common control comprising health care, financial, and consumer staples firms produces positive AI buildout exposure changes for all baseline-exposed sectors, with point estimates generally similar in magnitude to those in the baseline. Seven sectors retain FDR-adjusted significance, although only four also satisfy both parallel-pretrend tests, consistent with the common control providing a less tailored counterfactual than the portfolio-specific synthetic controls. These checks indicate that the direction and economic magnitude of the estimated AI buildout exposures are relatively insensitive to the factor and control specifications, even as statistical significance varies. The private-book figures reported in Section 6 are likewise insensitive to the treatment of unexited holdings whose most recent report predates 2025 and whose exit status is therefore uncertain, which we drop in the baseline; the alternatives, which include projecting exit dates with a predictive model or retaining all such positions as active, produce book-level figures that differ by small amounts.

Event days: the response to capital expenditure news

Table 4 examines whether the AI buildout and AI demand portfolios react differently to major AI-related announcements. We consider six events from July 2024 through February 2025, including hyperscaler capital-expenditure guidance, the Stargate announcement, and the DeepSeek shock. For each event, abnormal returns are measured using an ex-AI market model estimated over the prior 120 trading days, and we report the responses of the buildout and demand portfolios together with their difference.

The event responses differ markedly across the two portfolios. Capital-spending announcements consistently generate positive buildout-minus-demand abnormal returns, with especially large differentials following guidance from Microsoft, Alphabet, and Meta. The DeepSeek shock produces the opposite pattern, with a buildout-minus-demand abnormal return of -15.2 percent. The AI buildout portfolio declines 17.1 percent around the event, compared with a 1.9 percent decline for the AI demand portfolio, while several hyperscaler announcements are associated with positive buildout returns and comparatively weak demand responses. The evidence therefore suggests that the two portfolios capture different sensitivities to news about expected AI infrastructure investment rather than merely partitioning firms within a common AI category.

TABLE 4: Differential Reactions of AI Buildout and Demand to AI-Related Events

Event	Reaction Date	AI Buildout	AI Demand	AI Buildout – Demand
Microsoft AI Capex Guidance	7/31/2024	10.0%	0.3%	9.7%
Stargate Announcement	1/22/2025	3.5%	2.7%	0.8%
DeepSeek Shock	1/27/2025	-17.1%	-1.9%	-15.2%
Meta AI Capex Guidance	1/30/2025	0.7%	-3.3%	3.9%
Alphabet AI Capex Guidance	2/5/2025	2.8%	-2.4%	5.2%
Amazon AI Capex Guidance	2/7/2025	0.6%	-1.5%	2.2%

This table reports abnormal returns for the AI buildout and AI demand portfolios surrounding selected AI-related announcements. Abnormal returns are measured relative to a market model estimated over the 120 trading days before each event, with the ex-AI market excess return serving as the benchmark. The reaction date corresponds to the first trading day on which investors could respond to the news. The final column is the difference between the abnormal returns of the two portfolios. All returns are expressed in percentage terms.

6 The Private Book

The preceding section shows that the listed firms in eight real-asset sectors have acquired a sizable exposure to the AI buildout. Those estimates describe public markets, whereas private assets are not priced weekly, and their exposure cannot be estimated in the same manner. We instead ask what the public market estimates imply for a private portfolio that holds assets of the same types. We do so by taking the buildout exposure changes from the public portfolios and applying them to corresponding private market investments from closed-end funds provided by StepStone. This allows us to ask how much AI buildout exposure is already embedded in private real-asset portfolios, and whether that exposure reflects changes in the composition of new investment or rising sensitivity within assets already held.

The growing weight of new investment in power reflects a constraint, namely that the regulated grid cannot accommodate the power requirements of the buildout on the schedule that the buildout presumes. On the supply side, the generation and storage capacity waiting in U.S. interconnection queues now exceeds the entire installed fleet, and the median project takes five years from request to operation. On the demand side, the constraint is newer and steeper: the large-load requests before ERCOT, the grid operator for most of Texas, of which 90 percent are now data centers, have grown roughly sevenfold since the end of 2024, and in August 2026 ERCOT paused its large-load interconnection process pending a state-directed audit of every data center in its queue. A data center that cannot wait for a substation builds its own supply on its side of the

meter, through on-site generation financed by independent power producers, midstream firms, and private funds rather than through the rate base. Where the bottleneck is in the regulated system, the response is project-financed, and these constraints bear directly on the measurement of the private-capital mix-shift that follows.[16]

6.1 What private infrastructure holds

Deal count and infrastructure investment share by entry year for seven of the eight exposed sectors are shown in Figure 3; Building Products is omitted, as the sample contains no private investments in that sector over this period. The deal-level sample comprises 11,139 investments with entry dates from 2010 through 2025, representing $657 billion of invested capital.[17] The exposed sectors represent about 28 percent of invested capital in this deal-level sample. The figure is measured at cost; the corresponding share of NAV in the fund-level sample of Section 6.5 lies between roughly a third and two-fifths. The two figures differ in both basis and universe, and we state the basis on each occasion that an exposed share is reported. Within the exposed sleeve, independent power producers and renewables (IPP) account for 507 deals (41.7 percent of exposed-sector invested capital), followed by Midstream Energy with 253 deals (28.9 percent), Data Centers with 188 deals (17.3 percent), Electric Utilities with 67 deals (5.1 percent), Regulated Utilities ex Electric with 59 deals (3.9 percent), Construction & Engineering with 68 deals (2.9 percent), Heavy Electrical Equipment with 6 deals (0.2 percent), and Building Products with no deals in the sample.[18]

Data center deals rise from near zero before 2017 to 51 in 2025, becoming a substantial share of new infrastructure activity. IPPs and renewables account for a large, though volatile, share of infrastructure investment after 2020, while midstream's investment share has declined markedly from its earlier levels. Because investments enter the data with a reporting lag, activity in the most

[16] On the supply side, Lawrence Berkeley National Laboratory, *Queued Up: 2026 Edition* (June 2026), reports 2,061 GW of generation and storage capacity active in U.S. queues at end-2025 against an installed fleet of 1,374 GW, and a median of 61 months from request to operation for projects completed in 2025. On the demand side, ERCOT reported 63 GW of large-load requests in December 2024, 226 GW in December 2025, and approximately 474 GW as of June 2026, about 90 percent of it data centers (ERCOT, July 29, 2026); the pause followed the Governor's directive of August 3, 2026. In response, Enverus projects that roughly 36 percent of U.S. data center capacity added over 2026–2030 will be served behind the meter (August 2026).

[17] We drop 2,126 deals with no exit date and a report date before 2025, as their exit status in 2025 is uncertain. Subsequent results show little difference when these deals are handled differently, for example by predicting exit dates with a machine learning model or by assuming that all remain active through the end of the sample.

[18] There is some ambiguity as to whether data centers should qualify as real estate or infrastructure. For purposes of this paper, we define data centers as infrastructure.

recent years is mechanically understated; 2025 in particular should therefore be interpreted as a lower bound rather than a complete measure of activity.

EXPOSED-SECTOR ACTIVITY IN PRIVATE INFRASTRUCTURE

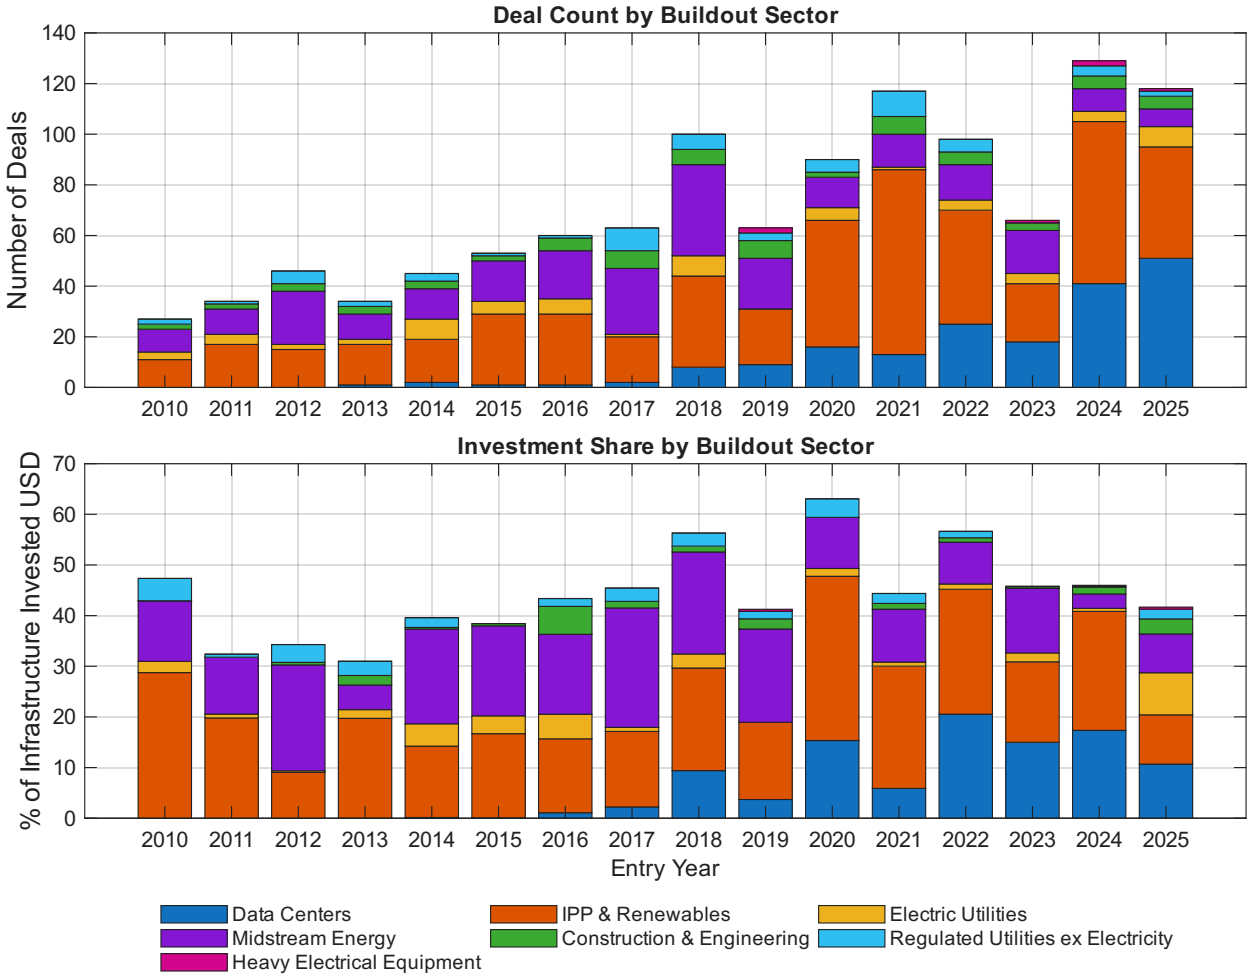


FIGURE 3: This figure shows annual investment activity across the seven exposed infrastructure sectors from 2010 through 2025. The top panel reports the number of new deals by entry year, while the bottom panel reports each sector's share of total infrastructure invested capital in that entry year. Bars are stacked by sector and ordered according to their 2025 investment share of infrastructure. Building Products is omitted because there are no private Building Products investments in the sample.

6.2 Mapping sector sensitivities to a book beta: stock and flow

To translate the public market results into private markets, we use the buildout exposure changes as measures of the change in buildout sensitivity associated with each real-asset sector. We apply these estimates to the corresponding private market portfolio weights in two ways. The stock measure weights each sector's exposure change by its share of total active synthetic NAV, cap-

turing the buildout exposure embedded in the existing portfolio, while the flow measure weights the same exposure changes by each sector's share of new invested capital, capturing the exposure associated with current investment activity.[19] The stock measure thus describes the portfolio that an investor owns today, and the flow measure describes the portfolio that the most recent year of investment is building. As an example of the calculation, a private book with 20 percent of its value in independent power producers, a sector for which the buildout exposure change is 0.664, acquires an implied buildout exposure of 0.13 from that position alone.

The book's implied change in buildout beta is the holdings-weighted sum of the sector changes, and its movement between a base year 0 and year t splits into a mix-shift and a risk-shift:

$$\begin{aligned} \Delta\beta^{B}_{\text{book},t} &= \sum_{s} w_{s,t}\,\Delta\beta^{B}_{s,t}, \\ \Delta\beta^{B}_{\text{book},t} - \Delta\beta^{B}_{\text{book},0} &= \underbrace{\sum_{s} \left(w_{s,t} - w_{s,0}\right)\overline{\Delta\beta^{B}_{s}}}_{\text{mix-shift}} + \underbrace{\sum_{s} \overline{w}_{s}\left(\Delta\beta^{B}_{s,t} - \Delta\beta^{B}_{s,0}\right)}_{\text{risk-shift}}. \end{aligned} \tag{9}$$

with weights $w_{s,t}$ either synthetic-NAV shares (stock) or vintage-year invested-capital shares (flow), bars denoting averages across the two dates, and the unclassified remainder of the book carrying $\Delta\beta^{B} = 0$ by assumption. This is the shift-share form of the holdings-implied loading in Dittmar and Lundblad (2017), applied to sectors rather than firm characteristics. The implied exposure of the book can increase for two reasons. The sectors that the book holds can become more sensitive to the buildout, which is the risk-shift, or the book can hold more of the sensitive sectors, which is the mix-shift. An investor whose sector weights remain unchanged would experience only the risk-shift.

6.3 The implied exposure of the book in 2025

Figure 4 reports the implied buildout exposure of the private book, computed as in equation (9) from the public market exposure changes of Table 3 and the sector weights of the StepStone holdings. We report 2025, the year in which the public comparables reprice. At year-end 2025, the

[19] In the absence of regularly reported NAVs, synthetic NAV estimates each investment's annual value by interpolating between invested capital at entry and its reported terminal value, or to its latest reported value if the investment remains active, using its reported total value multiple (TVM). Also note that some unexited holdings have reporting dates before 2025, making their active status in 2025 unknown. We drop these holdings from the main analysis, but the results are robust to other ways of handling them.

exposed sectors account for 35.4 percent of total active synthetic NAV in the deal-level sample, compared with 28.1 percent of cumulative invested capital across deals entered from 2010 through 2025. In earlier years, the implied exposure of the book is negligible, as the public market estimates on which it rests are themselves close to zero.

For the stock of private holdings, the implied buildout exposure change between 2019 and 2025 is 0.173, of which the risk-shift accounts for 0.157 and the mix-shift for 0.016. For the flow of new investment, the corresponding exposure change is 0.131, of which the risk-shift accounts for 0.116 and the mix-shift for 0.014 (the components do not sum exactly to the total on account of rounding). The magnitude should be read with the construction in mind. The portion of the book outside the eight exposed sectors is assigned no exposure change, and the exposed sectors represent roughly 28 percent of the book at cost (Section 6.1) and a somewhat larger share at NAV. A figure of 0.17 for the whole book therefore reflects sizable exposure changes within the exposed sectors, diluted by the portion of the book that those sectors do not represent. Were private assets priced at the frequency of their public comparables, a 10 percent decline in the buildout factor would be associated with roughly 1.7 percentage points of additional downside for the book as a whole, relative to the synthetic-control counterfactual, and a substantially larger incremental effect within the exposed sleeve.

The sector contributions are the product of weight and exposure change. Independent power producers and renewables contribute the most, as the sector combines the largest exposure change among the sectors that private funds hold in size with the largest weight in the infrastructure book; data centers and midstream energy follow. Electric utilities, a small share of the stock, contribute more to the flow, which reflects new investment in the sector in 2025.

The book-level figure is an average, and an allocator holds funds rather than the book. The next two subsections ask how the exposure is distributed across funds and across vintages, and whether it arrived through new investment or through positions already held. Section 6.6 then returns to the decomposition and to what its two components imply.

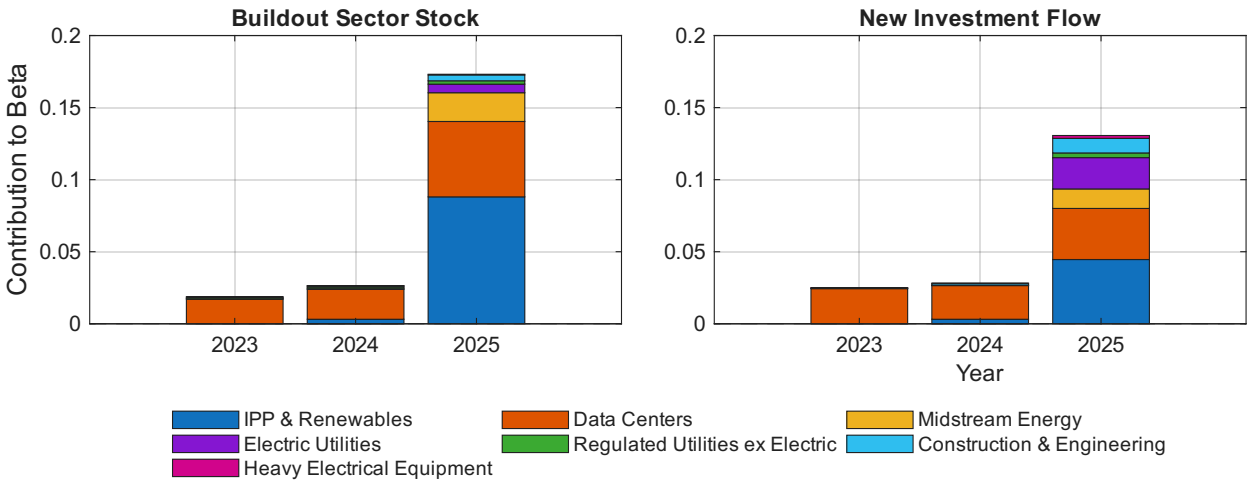


FIGURE 4: This figure decomposes the implied buildout exposure change of the private book into the contributions of the seven exposed sectors with private holdings, for 2023–2025. The left panel uses year-end synthetic NAV weights, with each sector's synthetic NAV scaled by total active real-asset synthetic NAV (the stock). The right panel uses new-investment weights, with invested capital in each sector's newly entered deals scaled by total real-asset investment entering in that year (the flow). Each sector's contribution is its portfolio weight multiplied by the sector's public market buildout exposure change from Table 3; the 2025 contribution uses the January 2025–March 2026 estimate. Bars are stacked by sector and ordered by 2025 stock weight. Building Products is omitted because there are no private Building Products investments in the sample. Between 2019 and 2025, the stock exposure change of 0.173 decomposes into a risk-shift of 0.157 and a mix-shift of 0.016. The corresponding flow exposure change of 0.131 decomposes into a risk-shift of 0.116 and a mix-shift of 0.014; components may not sum to totals because of rounding. The risk-shift weights changes in sector exposures by the average of 2019 and 2025 sector weights, while the mix-shift weights changes in sector weights by the average of sector exposures across the two dates.

6.4 The fund cross-section

A fund's implied buildout exposure change is the holdings-weighted sum of the exposure changes of the sectors it holds, computed as in equation (9) with the fund's actual weights. Because every private investment in a given sector is assigned the same public market estimate, dispersion across funds is dispersion in holdings. The cross-section therefore answers the question of whether the exposure is spread thinly across the fund universe or carried by a subset of funds.

The fund-level analysis requires complete fund holdings, and Section 3 describes the two samples we consider: the 159 clean funds for which every holding is observed, and the reconstructed panels that add, on average, roughly 335 funds whose missing holdings are filled in by the procedure described in Appendix E. The fund-level sample differs from the deal-level sample of Section 6.3 in one respect: rather than excluding investments entered before 2010, it excludes funds

whose entire observed portfolio had exited before 2010, such that a fund active after 2010 is observed with its full history. We report both samples throughout, and because clean funds are smaller and less exposed than the funds that require reconstruction, the clean-fund estimates are best read as a lower bound. We classify funds using their aggregate time-weighted synthetic NAV over 2023–2025, holding group membership fixed across years. A fund is *majority-exposed* when at least half of this NAV is in the eight exposed sectors, *minority-exposed* when the share is positive but below one-half, and unexposed otherwise.

Figure 5 reports the average fund-level exposure change by group. Before 2025, the averages are close to zero in every group, as the public market estimates are. In 2025, the average majority-exposed fund acquires a buildout exposure change ranging from 0.34 to 0.45, depending on the sample and whether funds are weighted equally or by NAV; the corresponding change for minority-exposed funds ranges from 0.07 to 0.09. Across all funds, including those with no exposure, the equally weighted averages are 0.075 in the clean sample and 0.111 across the reconstructed panels, rising to 0.118 and 0.186, respectively, when funds are weighted by NAV. The NAV-weighted average across all funds is of the same order as the book-level figure of Section 6.3, as it should be, with the difference between the two reflecting the different universes on which they are computed. In 2025, the clean-sample estimates are lower than those from the reconstructed panels for majority-exposed funds and for the full fund sample, but slightly higher for minority-exposed funds.

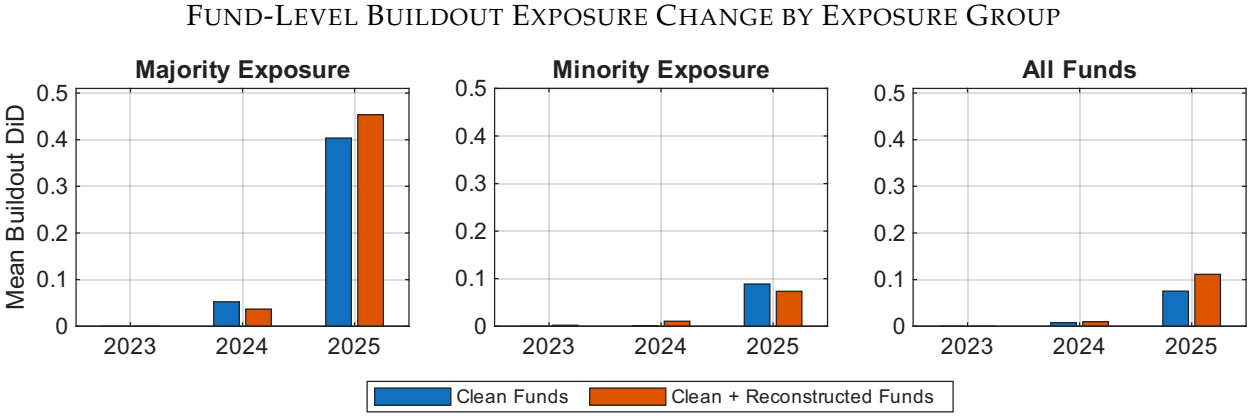


FIGURE 5: This figure reports the mean fund-level buildout exposure change by exposure group for 2023–2025. Majority-exposed funds hold at least 50% of NAV in the exposed sectors; minority-exposed funds hold a positive share below 50%. The clean-fund sample sizes are 18, 17, and 15 for majority exposure; 11, 8, and 9 for minority exposure; and 126, 116, and 91 for all funds in 2023, 2024, and 2025, respectively. The corresponding average sample sizes across the top 20 clean-plus-reconstructed panels are 77, 79, and 79; 62, 64, and 65; and 364, 368, and 364 funds. Blue bars report estimates using clean funds only; orange bars report the mean estimate across the top 20 reconstructed panels.

The averages conceal an asymmetric distribution. Figure 6 reports the distribution of fund-level exposure changes in 2025, restricted to funds with at least two active investments so that the right tail is not an artifact of single-holding funds. Among majority-exposed funds, the mass lies at positive values, with a mode between 0.4 and 0.5 and a secondary concentration near 0.7, the exposure change of independent power producers, reflecting funds that hold power assets nearly exclusively, whereas the median real-asset fund has no exposure.

Figure 7 makes the same point in terms of contribution to the book. Across the 364 funds in an average reconstructed panel, about 62 percent contribute nothing to the book's buildout exposure change, and the ten largest contributors account for 45 percent of it. Those ten funds represent 20 percent of the book's NAV, such that their contribution runs at better than twice their weight. The concentration reflects intensity rather than size, as the most exposed funds are those whose holdings are most concentrated in power and data centers rather than the largest funds. Three nested fund groups appear in this section and the next: the 493 funds of Table 10 are those with positive cumulative invested capital through 2025; the 364 funds of Figure 7 are those with positive synthetic NAV and a valid buildout exposure change in 2025; and the 328 funds of Table 5 are those among them with a first observed investment in 2011 or later. The approximately 36 funds excluded from the vintage table have pre-2011 vintages.

For an allocator, the implication is that the book-level figure of Section 6.3 describes the real-asset sleeve as a whole and says little about any particular fund within it. A program whose managers specialize in power or data centers carries an implied exposure two to three times that of the book; a program weighted toward traditional real estate, transportation, or social infrastructure carries little or none. The fund level is also the level at which an investor can act, as commitments are made to funds and not to sectors.

DISTRIBUTION OF FUND-LEVEL EXPOSURE CHANGES IN 2025

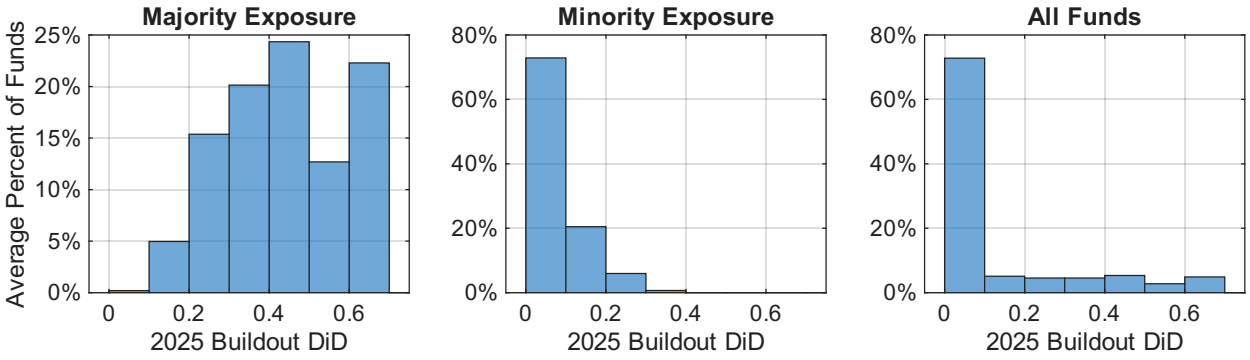


FIGURE 6: This figure shows the distribution of fund-level 2025 buildout exposure changes across the top 20 reconstructed panels. Fund-level exposure changes are computed with synthetic NAV weights across each fund's investments, and each histogram reports the average percentage of funds falling within each bin across the 20 panels. The sample is restricted to funds with at least two active investments in 2025. Majority exposure includes funds with at least 50% of NAV in the exposed sectors, minority exposure includes funds with a positive share below 50%, and all funds includes the full eligible sample. The corresponding average sample sizes across the 20 panels are 73, 65, and 333 funds, respectively.

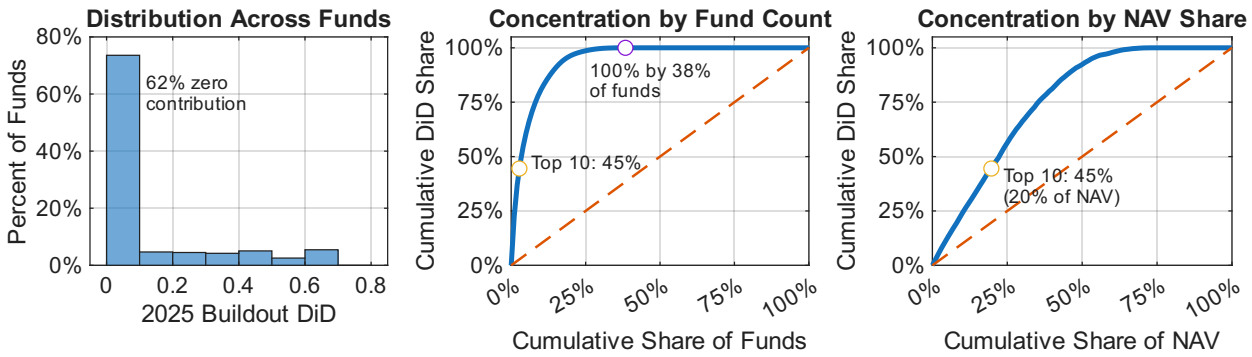


FIGURE 7: This figure shows the distribution and concentration of fund-level 2025 buildout exposure changes across the top 20 reconstructed panels, averaging 364 funds per panel. The left panel reports the average percentage of funds in each bin; approximately 62% of funds contribute nothing to the book-level exposure change. The middle panel ranks funds by their synthetic-NAV-weighted contribution to the book-level exposure change and plots the cumulative share of that total against the cumulative share of funds. The top 10 funds account for approximately 45% of the book-level exposure change, and the contributing funds account for approximately 38% of all funds. The right panel uses the same ranking but plots cumulative NAV share on the horizontal axis; the top 10 contributing funds represent approximately 20% of book NAV. Dashed 45-degree lines indicate equal contribution across funds or NAV, respectively.

6.5 Vintage and stock versus flow at the fund level

The fund cross-section establishes where the exposure sits. This subsection asks when it arrived. If the exposure of the private book were the product of recent investment in AI-related assets, it would be concentrated in the newest funds, and it would have grown within funds as capital was deployed. If it instead arrived through the repricing of positions already held, it would appear in old and new vintages alike, and it would be stable within funds over time. The second description is the one that the data support.

Table 5 reports 2025 fund-level exposure changes, and the composition of the exposed holdings, by fund vintage, with vintage defined by the fund's earliest observed investment. The exposure change does not increase systematically with vintage. Among majority-exposed funds, which represent roughly a fifth to a quarter of the funds in each vintage bucket, the average lies between 0.4 and 0.5 in every bucket, and funds with first observed investments in 2011–2013 carry exposure changes comparable to those entering a decade later. Across all funds, the NAV-weighted average lies between roughly 0.16 and 0.20 in every bucket, and the median is approximately zero in every bucket. The pattern suggests that exposure arose through positions already held, rather

than exclusively through new commitments.

What does vary with vintage is the composition of the exposed holdings. For the oldest vintages, the exposed sleeve consists of independent power producers, midstream energy, and regulated utilities, which is the legacy infrastructure portfolio. Data centers account for just 0.4 percent of NAV in the oldest vintage, compared with 14.6 percent in the newest, where they represent roughly a third of exposed holdings. The newest vintages are therefore only modestly more exposed than the oldest, but they are exposed through different assets, and through assets whose contract terms, counterparties, and residual value differ from those of a pipeline or a regulated utility, a distinction to which we return in Section 7.

TABLE 5: Fund-Level Buildout Exposure Change and Composition in 2025, by Fund Vintage

Panel A: Majority-Exposed Funds

Vintages	**2011–2013**	**2014–2016**	**2017–2019**	**2020–2022**	**2023–2025**
Number of funds	9	15	18	18	12
Mean exposure change	0.397	0.472	0.404	0.473	0.447
Median exposure change	0.384	0.496	0.370	0.478	0.437
NAV-weighted exposure change	0.451	0.441	0.381	0.461	0.364
Share of NAV in exposed sectors	93.4%	87.6%	78.2%	83.1%	77.9%
IPP & Renewables	50.0%	48.6%	31.4%	43.6%	24.4%
Data Centers	0.3%	6.8%	20.2%	29.0%	27.5%
Midstream Energy	26.0%	26.3%	19.2%	7.2%	23.0%
Electric Utilities	6.7%	3.2%	2.3%	1.1%	2.1%
Regulated Utilities ex Electric	10.1%	2.6%	3.1%	0.6%	0.7%
Construction & Engineering	0.3%	0.1%	1.9%	1.7%	0.3%

Panel B: Minority-Exposed Funds

Vintages	**2011–2013**	**2014–2016**	**2017–2019**	**2020–2022**	**2023–2025**
Number of funds	3	11	15	22	12
Mean exposure change	0.094	0.047	0.060	0.079	0.109
Median exposure change	0.088	0.026	0.045	0.067	0.088
NAV-weighted exposure change	0.077	0.040	0.061	0.086	0.101
Share of NAV in exposed sectors	19.2%	9.7%	17.5%	18.2%	19.4%
IPP & Renewables	6.1%	2.0%	1.4%	3.3%	9.0%
Data Centers	0.9%	2.4%	4.3%	10.1%	5.0%
Midstream Energy	10.4%	3.8%	10.7%	2.5%	3.4%
Electric Utilities	0.2%	1.2%	0.2%	0.1%	1.3%
Regulated Utilities ex Electric	1.1%	0.2%	0.6%	0.6%	0.0%
Construction & Engineering	0.5%	0.1%	0.3%	1.5%	0.6%

Panel C: All Funds

Vintages	**2011–2013**	**2014–2016**	**2017–2019**	**2020–2022**	**2023–2025**
Number of funds	36	62	82	96	52
Mean exposure change	0.110	0.121	0.099	0.107	0.125
Median exposure change	0.000	0.000	0.000	0.000	0.000
NAV-weighted exposure change	0.163	0.185	0.176	0.197	0.203
Share of NAV in exposed sectors	33.7%	37.3%	37.6%	36.6%	42.7%
IPP & Renewables	18.1%	19.4%	13.3%	16.7%	14.2%
Data Centers	0.4%	3.7%	9.8%	14.3%	14.6%
Midstream Energy	10.4%	11.4%	11.3%	3.6%	11.8%
Electric Utilities	1.5%	1.6%	1.0%	0.5%	1.4%
Regulated Utilities ex Electric	3.1%	1.1%	1.4%	0.4%	0.3%
Construction & Engineering	0.1%	0.1%	0.9%	1.2%	0.3%

This table reports fund counts, the mean, median, and NAV-weighted fund-level buildout exposure change in 2025, and the share of 2025 synthetic NAV held in each exposed sector, by fund-vintage bucket. Vintage is the year of the fund's earliest observed investment. Results are shown separately for majority-exposed funds (at least 50% of NAV in the exposed sectors), minority-exposed funds (a positive share below 50%), and all funds. Fund counts and statistics are averaged across the top 20 reconstructed panels. Building Products and Heavy Electrical Equipment are included in the exposed share but represent negligible shares of aggregate 2025 synthetic NAV and are omitted from the sector rows.

The same conclusion follows from tracing a fixed set of funds through time. Table 6 sorts the

funds active before 2023 into three groups by the share of their 2021–2022 NAV in the exposed sectors and holds group membership fixed thereafter. The three groups remain broadly where they began. The low-exposure group moves from 2 to 5 percent of NAV in the exposed sectors by 2025, the medium group from roughly 50 to 57 percent, and the high group remains near 90 percent. Table 7 compares funds that first invested in 2021 or earlier with those that first invested later. The legacy funds hold a steady share in the mid-to-high 30s throughout, while the newer funds begin with lower exposure shares and exhibit shares comparable to those of legacy funds in the later comparisons. Neither comparison suggests a broad reallocation of the existing fund universe toward the exposed sectors. Part of the persistence is mechanical, as a closed-end fund that has deployed its capital does not turn over its portfolio, and part reflects the lag with which recent investment enters the data, such that the absence of a response in the newest cohorts is provisional. It is nonetheless the pattern produced by a risk-shift (as opposed to a mix-shift).

TABLE 6: Share of NAV in the Exposed Sectors, by Pre-2023 Exposure Group

2021–2022 exposure group	**Low**	**Medium**	**High**
Number of funds	218	24	46
Exposed share of NAV, 2021–2022	1.9%	51.0%	92.9%
Exposed share of NAV, 2023	4.3%	53.3%	91.8%
Exposed share of NAV, 2024	4.5%	55.4%	90.2%
Exposed share of NAV, 2025	4.5%	57.2%	88.6%

The table reports the share of aggregate synthetic NAV held in the exposed sectors, over time, for a fixed set of funds with positive synthetic NAV in each year from 2021 through 2025. Funds are assigned to groups by the share of their aggregate 2021–2022 synthetic NAV in the exposed sectors: low (up to one-third), medium (above one-third through two-thirds), and high (above two-thirds). Group membership is held fixed thereafter. Fund counts and shares are averaged across the top 20 reconstructed panels.

TABLE 7: Share of NAV in the Exposed Sectors over Adjacent Years: Legacy and Post-2021 Funds

Adjacent years	**2022–2023**	**2023–2024**	**2024–2025**
Legacy funds (number)	317	302	288
Exposed share, start of period	35%	35%	36%
Exposed share, end of period	35%	36%	37%
Post-2021 funds (number)	23	47	65
Exposed share, start of period	28%	36%	38%
Exposed share, end of period	33%	38%	38%

This table compares the share of aggregate synthetic NAV held in the exposed sectors over adjacent years for legacy and post-2021 funds. Legacy funds are those whose first observed investment occurred in 2021 or earlier. Post-2021 funds are those whose first observed investment occurred in 2022 or later and no later than the first year of each comparison; the eligibility window therefore expands across comparisons, from the 23 funds that first invested in 2022 to the 65 that first invested between 2022 and 2024. Each comparison includes only funds with positive synthetic NAV in both years. Fund counts and shares are averaged across the top 20 reconstructed panels.

Table 8 looks inside the exposed sleeve. Data center transactions have become the most numerous form of new exposed investment, and their share of the active exposed stock, measured by NAV, rises from roughly a sixth in 2022 to nearly a third in 2025, while the share of midstream energy declines as older positions are realized and few new ones are entered. Measured in dollars of new investment, however, data centers represented between a third and two-fifths of new exposed capital in each year from 2022 through 2024, and roughly a quarter in 2025, a year in which electric utilities accounted for an unusually large share of the new exposed dollars. The change reflects the accumulation of data center transactions and the exit of midstream positions rather than a sustained increase in data centers' share of new exposed capital, and independent power producers remain the largest holding by value even as their weight declines. The mix-shift observed in the holdings to date is thus a change in the character of the exposed sleeve, from contracted energy transport toward data centers, with power generation remaining a major component, rather than a change in its size, which is why the mix-shift term in Section 6.3 is small. Because investments enter the data with a lag of up to three years, the new-deal counts and sizes for the most recent years describe an early-reported subset, and the flow figures for 2024 and 2025 are lower bounds on activity.

TABLE 8: Composition of the Exposed Sleeve, 2022–2025

Panel A: New deal count	**2022**	**2023**	**2024**	**2025**
Independent Power Producers & Renewables	45	23	64	44
Heavy Electrical Equipment	0	1	2	1
Construction & Engineering	5	3	5	5
Electric Utilities	4	4	4	8
Data Centers	25	18	41	51
Midstream Energy	15	17	9	7
Regulated Utilities ex Electric	5	0	4	2

Panel B: Share of new exposed capital	**2022**	**2023**	**2024**	**2025**
Independent Power Producers & Renewables	43%	35%	51%	23%
Heavy Electrical Equipment	0%	0%	0%	1%
Construction & Engineering	2%	1%	3%	7%
Electric Utilities	2%	4%	1%	20%
Data Centers	36%	33%	38%	26%
Midstream Energy	16%	28%	6%	18%
Regulated Utilities ex Electric	2%	0%	0%	5%

Panel C: Active deal count	**2022**	**2023**	**2024**	**2025**
Independent Power Producers & Renewables	223	223	270	285
Heavy Electrical Equipment	2	2	4	4
Construction & Engineering	30	30	34	36
Electric Utilities	18	18	21	25
Data Centers	68	84	120	148
Midstream Energy	135	136	128	102
Regulated Utilities ex Electric	37	35	35	35

Panel D: Share of active exposed NAV	**2022**	**2023**	**2024**	**2025**
Independent Power Producers & Renewables	44%	42%	42%	38%
Heavy Electrical Equipment	0%	0%	0%	0%
Construction & Engineering	2%	2%	2%	2%
Electric Utilities	3%	3%	3%	4%
Data Centers	17%	22%	27%	31%
Midstream Energy	29%	28%	24%	22%
Regulated Utilities ex Electric	4%	4%	3%	3%

The table reports the composition of private holdings in the exposed sectors from 2022 through 2025. New deal count and the share of new exposed capital describe investments entered in each year; active deal count and the share of active exposed NAV describe the stock of active investments in the exposed sectors in each year. Shares are of the exposed sleeve, not of the book.

6.6 An established risk-shift and an emerging mix-shift

The decomposition of Section 6.3 attributes the larger share of the 2025 increase to the risk-shift, and the fund-level evidence of Sections 6.4 and 6.5 is consistent with that attribution: the expo-

sure is carried by funds of every vintage, it is stable within funds, and the movement within the exposed sleeve is a rotation among exposed sectors rather than growth in the sleeve. We regard the two components as differing in kind.

The risk-shift is established, in that it is the product of two quantities that are each well measured: the change in the buildout sensitivity of the public comparables, for which inference is reported in Table 3, and the sector weights of the existing private book, which is dominated by positions entered in earlier years for which reporting is largely complete. The risk-shift, moreover, requires no action on the part of the investor. A private fund that has held the same power and midstream assets since 2015 acquired the exposure by virtue of the repricing of those sectors in public markets.

The mix-shift, in contrast, is modest to date, and we regard the figures as understating the contribution that will eventually be observed, for three reasons. The first is the reporting lag. Private investments entered in the most recent years are only partially observed, and the flow weights for those years rest on an incomplete sample. The second is the timing of deployment. Capital committed to a private fund is drawn over an investment period of several years, and a commitment made in 2025 enters a holdings-based measure only as the capital is called. A sizable portion of the capital raised for data center and power strategies in recent years has, by construction, yet to reach the holdings data. The third is the form in which the exposure is acquired. Exposure obtained through co-investments, direct positions, separately managed accounts, and open-end vehicles is recorded with delay, and sometimes not at all, since the source database is organized primarily around closed-end drawdown funds.

The two components also compound. Under the symmetric decomposition, the mix-shift weights changes in sector weights by the average of the sector exposure changes across the two dates, thereby allocating the interaction between changing weights and changing sensitivities equally across the risk-shift and mix-shift components. Capital that enters independent power producers or data centers in 2026 is applied to a buildout coefficient of roughly 0.7 or 0.5, whereas the same capital entering in 2023 would have been applied to a coefficient near zero. A mix-shift of a given size therefore adds more to the exposure of the book than it would have before the repricing.

6.6.1 The direction of capital not yet deployed

The holdings evidence of Section 6.5 and the fundraising evidence of this subsection are consistent, and the reconciliation is the distinction between capital deployed and capital committed. The existing fund universe is not reallocating: a closed-end fund that has deployed its capital does not turn over its portfolio (although it may have the opportunity for limited add-on investments or recycling), and the exposed share of its NAV accordingly moves little. The tilt toward the exposed sectors appears instead in the capital that has yet to be deployed, the dry powder of existing funds and the capital being raised for new ones, which is absent from the holdings-based measure until it is called. The holdings data do not yet reveal the likely extent of the prospective mix-shift, but fundraising provides an indication of its direction. We therefore extend the analysis to capital being raised for 2026-vintage funds, most of which remain in fundraising or pre-marketing and have deployed little capital. For each general partner, we estimate the exposed-sector share of its reconstructed 2023–2025 holdings and apply that share to the capital associated with its new funds, effectively assuming that new fundraising reflects the manager's recent investment mix; Appendix G provides additional details. This measure captures capital raised by managers with holdings in the exposed sectors rather than capital explicitly raised for AI infrastructure, and the 2026 estimates annualize reported fund sizes and targets through the third quarter.[20]

As Table 9 reports, estimated capital attributable to the exposed sectors represents 36.5 percent of real-asset fund capital for the 2023 vintage, falls to 21.0 percent for 2024, a weak year for fundraising by exposed managers, and recovers to 37.8 percent for 2025 and 53.7 percent for 2026. The 2026 share does not depend on the annualization, which scales the exposed and other components alike, but the dollar level does: reported fund capital for 2026-vintage funds is about $430 billion through the third quarter, against $397 billion for all of 2025, and the $573 billion total in the table annualizes that amount.

[20]The amounts are equity invested, or available for investment, by the closed-end funds in the StepStone sample, measured before fund-level and asset-level leverage. They exclude private credit, open-end vehicles, co-investments, and direct positions. The figures are accordingly small relative to the capital expenditure totals of Section 1, and the informative quantities are the shares rather than the dollar levels.

TABLE 9: Estimated Fund Capital Attributable to the Exposed Sectors, by Fund Vintage, 2023–2026

Fund capital, final size or target (USD billions)	2023	2024	2025	2026 (annualized)
Exposed sectors	163.8	72.3	150.0	307.5
Other sectors	285.2	271.8	247.3	265.2
Exposed share	36.5%	21.0%	37.8%	53.7%

This table reports estimated fund capital attributable to the exposed sectors and to the other sectors, in USD billions. Within each of the 20 reconstructed panels, each general partner's exposed-sector share of NAV over 2023–2025 is applied to the fund capital associated with that general partner, where fund capital is reported fund size for closed funds and target fund size for funds still fundraising or in pre-marketing. The resulting estimates are pooled, extrapolated to the full real-asset fund universe within each reconstruction using the corresponding fund-capital coverage adjustment, and averaged across reconstructions. The 2026 values annualize fund capital reported through the third quarter (about $430 billion) and should be interpreted as estimates rather than observed full-year fundraising. Fund capital is assigned to a fund's vintage year and includes the target size of funds that have yet to hold a final close; the figures are therefore not comparable to industry tallies of capital closed within a calendar year, which are lower, particularly for 2026, as most 2026-vintage funds remain in fundraising or pre-marketing.

The estimates are indicative rather than forecasts, and we note four qualifications. Fund targets are not closes; funds that fail to raise have historically represented about 6 percent of a vintage's fund capital, and at most 11 percent. For 2026, the allocation across individual exposed sectors is held at its 2025 composition, while the exposed share of overall fundraising is allowed to vary with the mix of managers raising capital. Using each manager's 2023–2025 holdings to infer the exposed-sector share of new fundraising understates the shift if managers increase their holdings in these sectors, particularly if previously unexposed managers enter them, and overstates it if exposed managers diversify away from them. Finally, the annualization of three quarters assumes that fundraising in the final quarter proceeds at a similar pace.

Dry powder provides an additional source of capital available for future deployment. We estimate year-end 2025 dry powder as each fund's remaining deployable capital after accounting for cumulative deployment, assuming that 90 percent of reported fund size is available for investment. We allocate this capital between the exposed and the other sectors according to each fund's existing investment mix and average the results across the 20 reconstructed fund panels. Further details on the methodology are provided in Appendix G.

The results are shown in Table 10. The exposed sectors represent 30.0 percent of the capital that the funds in the sample have deployed through 2025, measured at cost. The exposed sectors account for 37.2 percent of time-weighted active synthetic NAV during 2025 and an estimated 45.5

percent of remaining dry powder at year-end 2025, a difference of 8.3 percentage points. Applied to the $255.1 billion of remaining deployable capital, this difference corresponds to roughly $21.1 billion more in the exposed sectors than the active-NAV mix would imply.

TABLE 10: Deployed Capital, Active NAV, and Remaining Dry Powder, by Sector Group

	Deployed Capital			Dry Powder
Fund Capital	**Through 2023**	**Through 2024**	**Through 2025**	**Year-End 2025**
Fund count	465	482	493	493
Total ($bn)	589.4	635.4	682.4	255.1
Exposed sectors ($bn)	178.3	191.3	205.0	*116.1*
Other sectors ($bn)	411.1	444.1	477.4	*139.0*
Exposed share	30.2%	30.1%	30.0%	*45.5%*
Other share	69.8%	69.9%	70.0%	*54.5%*
Exposed share of active NAV	*35.5%*	*36.5%*	*37.2%*	
Dry powder share of fund size	*34.1%*	*31.9%*	*29.7%*	*29.7%*
Difference vs. 2025 active NAV				*8.3 pp*

This table reports cumulative deployed capital, estimated active NAV, and remaining dry powder. Deployed capital includes all investments made through the end of each respective year, including exited investments, measured at cost. Active NAV includes investments with positive time-weighted synthetic NAV during each calendar year, with exposed-sector shares calculated using those weights. We assume that 90 percent of reported fund size is available for investment, with the remaining 10 percent reserved for fees and expenses. Estimated dry powder equals 90 percent of fund size less cumulative invested capital, with negative values set to zero. The exposed-sector allocation of remaining dry powder is estimated using each fund's cumulative exposed-sector share of invested capital. The difference in exposed share compares year-end 2025 dry powder (45.5 percent) with time-weighted active NAV during 2025 (37.2 percent), yielding 8.3 percentage points. Italicized numbers are estimates rather than directly observed values. All results are averages across the 20 reconstructed fund panels.

Combining the two pools, and assuming that 90 percent of new-vintage fund capital is available for investment, roughly half of the capital available to finance future private real-asset investment, $393 billion of $771 billion, is associated with the exposed sectors, against 37.2 percent of time-weighted active NAV during 2025 and 30.0 percent of the capital deployed to date. These figures describe a stock of capital available for eventual deployment rather than investment expected to occur during 2026, as private fund capital is deployed over several years.

Table 11 translates the two pools into an illustrative figure for 2026 deployment. It applies the 2023–2025 average rates at which beginning-of-year dry powder (11.5 percent) and new-vintage capital (17.8 percent) were deployed within the year to the 2026 pools, and scales the result to the deal-level coverage of the sample. Appendix G gives the steps and the corresponding figures on the broader fund-capital basis. On this illustration, deployment in the exposed sectors would rise

from roughly $15 billion a year over 2023–2025 to $34.5 billion in 2026, or from about 30 percent to 52 percent of real-asset deployment, with data centers and independent power producers accounting for most of the increase. Most of the projected increase follows from the annualized 2026 fundraising figure (subject to the qualifications noted above). The 52 percent figure is best read as the upper end of a range. The projection scales fund-level deployment to the deal-level sample with a single pooled factor of 0.540, whereas over 2023–2025 the ratio of observed to implied deployment was lower for the exposed sectors (0.433) than for the other sectors (0.60), such that the same procedure implies a higher exposed share of deployment than the deal-level data record in each of those years (Appendix Table A11). Scaling the two groups by their own ratios places 2026 deployment in the exposed sectors at roughly $28 billion, or about 44 percent of the total. The final row applies the public market exposure changes to each year's deployment-weighted sector composition, yielding an implied buildout exposure change that rises from zero in 2023 to 0.130 in 2025. Under the illustrative 2026 deployment mix, this increases to 0.253, or to roughly 0.21 at the lower end of the range, conditional on the latest estimated sector sensitivities persisting rather than representing an estimate of realized 2026 exposure.

TABLE 11: Illustrative 2026 Deployment in Private Real Assets, Deal-Level Basis

Deployment (USD billions)	**2023**	**2024**	**2025**	**2026**
Total real assets	46.8	50.8	53.8	*66.6*
Exposed sectors	15.2	14.2	15.5	*34.5*
IPP & Renewables	5.3	7.3	3.6	*12.9*
Data Centers	5.0	5.4	4.0	*10.6*
Midstream Energy	4.2	0.9	2.8	*7.6*
Other Exposed Sectors	0.7	0.6	5.1	*3.4*
Exposed share	32%	28%	29%	*52%*
Implied buildout exposure change	0.000	0.027	0.130	*0.253*

Deployment is reported in USD billions. The 2023–2025 columns report observed deal-level deployment. The 2026 column (italics) is illustrative: the 2023–2025 average within-year deployment rates are applied to year-end 2025 dry powder and annualized 2026-vintage fund capital, and the result is scaled to the deal-level coverage of the sample using a common pooled 2023–2025 deal-to-fund scaling factor of 0.540. Applying separate scaling factors to the exposed sectors (0.433) and to the other sectors (0.60) in place of the pooled factor would give 2026 exposed-sector deployment of roughly $28 billion, an exposed share of about 44 percent, and an implied buildout exposure change of roughly 0.21; the 2026 column is accordingly the upper end of a range. Other Exposed Sectors represent the difference between total exposed-sector deployment and the three separately reported subsectors. The implied buildout exposure change is the deployment-weighted average of the public-market synthetic-control exposure changes, using the year-specific estimates for 2023 and 2024 and the January 2025–March 2026 estimates for 2025 and 2026. The composition of Other Exposed Sectors is estimated from reconstructed private holdings. The 2026 implied exposure change is a conditional illustration rather than an estimate of realized exposure: it measures the exposure that projected deployment would carry if the latest public-market sector sensitivities persist. Appendix G describes the deployment projection methodology and reports the corresponding figures on the fund-capital basis.

While the risk-shift remains the dominant component, the projected 2026 deployment mix implies a substantially greater contribution from the mix-shift. At the upper end of the range, and relative to the 2019 baseline, the mix-shift accounts for approximately 28 percent of the implied buildout exposure of projected 2026 investment, compared with 11 percent for the flow of new investment in 2025 and 9 percent for the existing stock (Section 6.3). This suggests that new investment, if deployed as projected, would increasingly amplify the buildout risk already inherited through existing holdings.

We conclude that the evidence points to a prospective shift in the composition of private real-asset investment toward the exposed sectors. The precise magnitude is necessarily uncertain because the 2026 estimates rely on reconstructed portfolios, fundraising extrapolations, and

historical deployment relationships, but the direction of the result is consistent across the different measures: the exposed sectors represent a larger share of prospective capital and projected deployment than they have in recent years.

6.6.2 Evidence from capital expenditure forecasts

The private holdings data are backward-looking by construction. As a forward-looking check, we turn to I/B/E/S consensus forecasts of firm-level capital expenditure, which are available at several fiscal-year horizons and can be tracked as they are revised. We take two-year-ahead forecasts and measure the revision over a standardized six-month window, from the February consensus to the August consensus, in each year from 2022 through 2026, for the demand complex, the buildout tier, the eight exposed real-asset sectors, the ex-AI market, and a broader group of unexposed real-asset sectors. To ensure comparability across years and forecast horizons, we restrict each group to a constant set of firms with available forecasts in all five years, using the same firm sample as in Appendix Table A12. This coverage requirement retains 34 of the 58 firms classified as AI firms. Although the balanced-panel requirement reduces the AI sample, the retained firms account for 96.6 percent of the original AI universe's aggregate market capitalization as of 2026Q1.

Table 12 is read in three steps. The first step measures the shock. The firms of the demand complex are the firms doing the spending, and the revision to their expected capital expenditure, 22.8 percent over six months on average, is the buildout shock (the equipment suppliers of the buildout tier receive that spending as revenue, such that their own capital expenditure, revised by 7.8 percent, is not the relevant quantity). The second step is the response of the real-asset sectors. Expected capital expenditure in the eight exposed sectors is revised upward by 6.4 percent over the same windows, which is the counterpart, in fundamentals, of the repricing documented in Section 5: the listed owners of power, pipelines, and data center capacity are expected to build more. The third step is the placebo. Expected capital expenditure in the broader group of real-asset sectors with no direct link to the buildout is revised upward by only 0.6 percent on average, compared with 1.5 percent for the ex-AI market as a whole. The contrast between 6.4 and 0.6 percent is the forward-looking analogue of the exposed-versus-unexposed distinction in Table 3, although the unexposed comparison here includes additional real-asset sectors.

TABLE 12: Revisions in Two-Year-Ahead Capital Expenditure Forecasts, 2022–2026

		Average Capex Revision	
Group	**Firms**	**6-Month**	**Annualized**
All AI Firms	34	19.8%	47.1%
AI Demand Firms	11	22.8%	54.8%
AI Buildout Firms	20	7.8%	18.9%
Exposed Real-Asset Sectors	40	6.4%	13.4%
Ex-AI Real Assets	29	0.6%	1.3%
Ex-AI Market	627	1.5%	3.0%

This table reports revisions in I/B/E/S consensus capital expenditure forecasts. For each year from 2022 through 2026, we compare the February consensus forecast for capital expenditures two fiscal years ahead with the corresponding August forecast. The sample is restricted to firms with valid forecasts at both dates in all five years, as well as valid initial and latest forecasts for the corresponding fiscal years, ensuring a common balanced firm sample across this table and Appendix Table A12. The Firms column reports the number of firms in each group. Firm-level forecasts are aggregated within each group before calculating the six-month revision, and the resulting revisions are averaged across years. Annualized revisions compound each six-month revision over a twelve-month horizon before averaging across years. AI Demand Firms and AI Buildout Firms correspond to the demand complex and buildout tier in Table 1. Exposed Real-Asset Sectors are the eight sectors above the line in Table 3. Ex-AI Real Assets comprises Traditional REITs, Tower/Fiber REITs, other real estate firms, Integrated Oil & Gas, Airport Services, Highways & Railtracks, Rail Transportation, Marine Ports & Services, Marine Transportation, Passenger Ground Transportation, Environmental & Facilities Services, Diversified Support Services, and Data Processing & Outsourced Services, excluding all explicitly classified AI firms and the eight exposed real-asset sectors. Ex-AI Market is defined as in Section 2. Appendix G reports the underlying forecast levels by fiscal year.

Two implications follow. The revisions are consistent, from fundamentals rather than from prices, with the conclusion that the exposure of real assets runs through the buildout. They are also a further reason to regard the mix-shift of Section 6.3 as understated: the sectors in which private capital would have to be deployed are the sectors in which the listed firms are raising their investment plans most sharply.

Admittedly, a larger mix-shift is not a certainty. To the extent that investors respond to the risk-shift we document by limiting new commitments to the exposed sectors, the realized mix-shift could be smaller than the fundraising figures imply. Deployment into the exposed sectors is also subject to delivery risk, as the siting, permitting, and interconnection constraints described at the outset of Section 6 can delay or prevent the conversion of committed capital into operating assets. Eventually, the holdings data will permit a direct assessment as reporting for recent vintages is

completed. What can be said now is that the private book has already acquired exposure to the AI buildout through positions it has long held, that this exposure is concentrated in a minority of funds, and that new fundraising is increasingly associated with managers whose recent portfolios hold the exposed sectors in size. The next section considers how an investor might employ these estimates.

7 Implications for Allocators

The paper has operated from the perspective of an investor who manages a diversified portfolio and who has traditionally held real assets as diversifiers. We return to that perspective here. A natural objection to the exercise in Section 6 is that a beta is not the instrument with which a private book is managed: capital is committed years before it is drawn, positions cannot be easily adjusted at the margin, and the portfolio is governed by pacing plans and manager selection rather than by factor budgets. The objection has merit for the management of the real-asset program in isolation, but it carries considerably less force for the management of the total portfolio.

For example, consider a total portfolio approach in which the relevant quantity is the exposure of the fund to a given source of risk across every asset class in which that risk is held. A sleeve-level beta is of limited use in managing the sleeve itself. It is, however, exactly the input that the real-asset program must consider when making a fund-wide accounting of exposure to the AI capital expenditure program. Likewise, it is what is required for an assessment of whether real assets continue to diversify the remainder of the portfolio.

Importantly, AI buildout exposure is almost certainly held across the fund in three distinct payoff profiles. Listed equity holds the AI demand complex and the equipment suppliers. For public equities, the claim is convex in the commercial success of the technology. Real assets hold the physical capacity, generation, transmission, pipelines, and data centers, where the claim is an (often contracted) payment stream with a residual value exposed to obsolescence and, upon contract expiry, to merchant pricing. Credit holds the financing; the upside is capped, and losses materialize with a retreat in capital expenditure or with the obsolescence of the collateral. The three profiles correspond to the separation that Van Nieuwerburgh (2026) describes between the users of computing capacity, the owners of the underlying assets, and the investors who bear the financial risk. The vulnerabilities he identifies at the asset level, namely tenant concentration,

technological obsolescence, execution delay, and uncertain residual value, are borne in the first instance by the real-asset and credit sleeves rather than by the listed equity of the firms doing the spending.

The credit exposure has only recently begun to receive attention (Van Nieuwerburgh, 2026; Aldasoro et al., 2026).[21] AI capital expenditure is increasingly debt-financed at the margin, through investment-grade issuance by the hyperscalers, off-balance-sheet vehicles financed by private credit, securitizations backed by data center leases, project finance for new generation, and loans secured by the processors themselves. A fund can hold the hyperscalers in its public equity portfolio, their bonds in its credit portfolio, a data center platform in its infrastructure portfolio, and a direct loan against the same campus in its private credit portfolio. Two features of the financing make the exposure larger and less visible than the fund's own books would indicate. Leverage at the asset level can be far higher than the reported leverage of the tenant, as the Hyperion vehicle illustrates. A sizable portion of the hyperscalers' commitments takes the form of future leases and residual-value guarantees that are not yet recognized on their balance sheets, which Van Nieuwerburgh (2026) places at roughly $660 billion of $970 billion in lease commitments. An allocator who measures only the recognized debt of the tenant, or only the sector weights of the real-asset sleeve, will understate the claim that the total portfolio holds on the same assets and the same counterparties.

To be sure, these varied exposures are not necessarily additive across asset classes. An efficiency shock of the kind priced on January 27, 2025 reduces the value of capacity and of the equipment that fills it, whereas its effect on the platforms that purchase the capacity is ambiguous. Similarly, a retreat in capital expenditure may be welcomed by the shareholders of the firm doing the spending, yet it is clearly adverse for the owner of capacity built in anticipation of that spending and for the lender that financed it. It is the separation of the buildout factor from the demand factor in Section 2 that allows these offsetting effects to be measured. A single AI factor would, by construction, net them against one another.

Several practical implications follow from our analysis. First, institutional investors can report an implied buildout exposure for the private real-asset book, together with its decomposition

[21] We are not aware of any measurement of the corresponding fixed-income factor exposure. However, our measurement specification extends naturally to credit returns (or any other affected asset class), and we view this as an important direction for subsequent work.

into risk-shift and mix-shift. The calculation requires only sector-level holdings weights and the public market coefficients, and it can be performed for any real-asset portfolio. Under a total portfolio approach, the resulting figure can then be aggregated, with due attention to sign, with the corresponding exposures in public equity and credit.

Second, investors can monitor the mix-shift before it appears in holdings. The risk-shift is already present in the portfolio, whereas the mix-shift will be determined by capital that has been committed and has yet to be drawn. Applying the public market coefficients to unfunded commitments by target sector provides a forward-looking measure of the implied exposure, and it identifies the point at which pacing decisions can still affect the outcome. The variable to monitor on the other side of the ledger is the pace of capital expenditure commitments rather than the adoption of AI products, and the consensus forecast revisions of Section 6.6.2 are a readily available series for that purpose: a retreat in expected hyperscaler capital expenditure would be the first signal that the exposure documented here is being repriced.

Finally, a caveat applies to the listed proxy with which most risk systems represent private infrastructure and real estate: the relevant question for the risk team is what that proxy now holds, given that a proxy selected in 2019 represents a materially different portfolio in 2026.

8 Conclusion

Real assets are held as diversifying, long-duration investments with potential inflation protection. The types of assets held by private real-asset funds are also, to a considerable degree, the assets with which the AI capital expenditure program is being built, and public markets have begun to price them accordingly. In public markets, we document that several listed real-asset sectors acquire economically large and statistically significant buildout exposure, with no corresponding exposure to AI demand.

For private real assets, which are not valued at a comparable frequency, we apply the public market estimates to the sector composition of private holdings. The implied buildout beta of the private book increases in 2025, largely on account of a risk-shift, with a smaller role for a mix-shift. The risk-shift has already occurred, as it follows from the repricing of sectors that private funds already held. The mix-shift is modest to date, but reporting lags, the multi-year drawdown of recent commitments, and the growth of private data center and power investment lead us to

expect a larger contribution, applied to sensitivities that the risk-shift has already raised. The direction is already visible in the capital that has yet to be deployed: the dry powder of existing funds and the capital being raised for new ones are tilted toward the exposed sectors by a wider margin than the capital invested to date, and consensus forecasts of capital expenditure are being revised upward for the listed owners of the same types of assets. Our results indicate that the diversification premise on which real-asset allocations rest should be re-examined at the level of the total portfolio, where the exposure of real assets to the AI buildout is held alongside the related exposures of public equity and credit.

References

Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association*, 105(490):493–505, 2010.

Daron Acemoglu. The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58, 2025.

Iñaki Aldasoro, Sebastian Doerr, and Daniel Rees. Financing the AI boom: From cash flows to debt. BIS Bulletin 120, Bank for International Settlements, January 2026.

Aligned Data Centers. AI Infrastructure Partnership, MGX, and BlackRock's Global Infrastructure Partners to acquire all equity in Aligned Data Centers. Press release, October 15, 2025.

Fernando E. Alvarez, David Argente, Joyce Chow, and Diana Van Patten. Data centers and local economies in the age of AI: A shift–share approach. NBER Working Paper 35194, National Bureau of Economic Research, May 2026.

Aleksandar Andonov, Roman Kräussl, and Joshua D. Rauh. Institutional investors and infrastructure investing. *Review of Financial Studies*, 34(8):3880–3934, 2021.

Andrew Ang. *Asset Management: A Systematic Approach to Factor Investing*. Oxford University Press, New York, 2014.

Tania Babina, Anastassia Fedyk, Alex He, and James Hodson. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745, 2024.

Bank of England. Financial stability report. Technical report, Bank of England, Financial Policy Committee, December 2025.

Geert Bekaert and Xiaozheng Wang. Inflation risk and the inflation risk premium. *Economic Policy*, 25(64):755–806, 2010.

Jonathan B. Berk, Richard C. Green, and Vasant Naik. Optimal investment, growth options, and security returns. *Journal of Finance*, 54(5):1553–1607, 1999.

Robert J. Bianchi, Graham Bornholt, Michael E. Drew, and Michael F. Howard. Long-term U.S. infrastructure returns and portfolio selection. *Journal of Banking & Finance*, 42:314–325, 2014.

Ron Bird, Harry Liem, and Susan Thorp. Infrastructure: Real assets and real returns. *European Financial Management*, 20(4):802–824, 2014.

Frédéric Blanc-Brude, Timothy Whittaker, and Simon Wilde. Searching for a listed infrastructure asset class using mean–variance spanning. *Financial Markets and Portfolio Management*, 31(2): 137–179, 2017.

Marshall E. Blume. On the assessment of risk. *Journal of Finance*, 26(1):1–10, 1971.

Gregory W. Brown, Christian Lundblad, and William M. Volckmann II. Inflation hedging and real assets: Are public and private investments the same? *Institute for Private Capital working paper*, 2025. URL `https://uncipc.com/wp-content/uploads/2025/11/Inflation_and_Private_Assets_June-2025.pdf`.

Gregory W. Brown, Christian Lundblad, and William M. Volckmann II. Risk-adjusted performance of private funds: What do we know? *Institute for Private Capital working paper*, 2026. URL `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7413758`.

Ricardo Bustos and Andrew DeMond. The MSCI private infrastructure factor model. Research paper, MSCI Inc., September 2025.

Ricardo J. Caballero. Speculative growth and the AI "bubble". NBER Working Paper 34722, National Bureau of Economic Research, January 2026.

CAIA Association. Innovation unleashed: The rise of total portfolio approach. Technical report, Chartered Alternative Investment Analyst Association, March 2024.

CalPERS. CalPERS board adopts streamlined investment approach to seize market opportunities. Press release, November 17, 2025. URL `https://www.calpers.ca.gov/newsroom/calpers-news/2025/calpers-board-adopts-streamlined-investment-approach-to-seize-market-opportunities`.

Murray Carlson, Adlai Fisher, and Ron Giammarino. Corporate investment and asset price dynamics: Implications for the cross-section of returns. *Journal of Finance*, 59(6):2577–2603, 2004.

Spencer J. Couts, Andrei S. Gonçalves, and Andrea Rossi. Unsmoothing returns of illiquid funds. *Review of Financial Studies*, 37(7):2110–2155, 2024.

Robert F. Dittmar and Christian T. Lundblad. Firm characteristics, consumption risk, and firm-level risk exposures. *Journal of Financial Economics*, 125(2):326–343, 2017.

Andrea L. Eisfeldt, Gregor Schubert, Miao Ben Zhang, and Bledi Taska. Generative AI and firm values. *Journal of Finance*, 2026. doi: 10.1111/jofi.70089. Forthcoming.

Eugene F. Fama and Kenneth R. French. Industry costs of equity. *Journal of Financial Economics*, 43 (2):153–193, 1997.

Eugene F. Fama and G. William Schwert. Asset returns and inflation. *Journal of Financial Economics*, 5(2):115–146, 1977.

David Geltner. Estimating market values from appraised values without assuming an efficient market. *Journal of Real Estate Research*, 8(3):325–345, 1993.

David M. Geltner. Smoothing in appraisal-based returns. *Journal of Real Estate Finance and Economics*, 4(3):327–345, 1991.

Mila Getmansky, Andrew W. Lo, and Igor Makarov. An econometric model of serial correlation and illiquidity in hedge fund returns. *Journal of Financial Economics*, 74(3):529–609, 2004.

Robin Greenwood and Samuel G. Hanson. Waves in ship prices and investment. *Quarterly Journal of Economics*, 130(1):55–109, 2015.

David Hartzell, John S. Hekman, and Mike E. Miles. Real estate returns and inflation. *AREUEA Journal*, 15(1):617–637, 1987.

Martin Hoesli and Elias Oikarinen. Are REITs real estate? Evidence from international sector level data. *Journal of International Money and Finance*, 31(7):1823–1850, 2012.

Holly Horrigan, Brad Case, David Geltner, and Henry Pollakowski. REIT-based property return indices: A new way to track and trade commercial real estate. *Journal of Portfolio Management*, 35(5):80–91, 2009.

Wendy Hu, Christian T. Lundblad, and Vedant Mozumdar. The investment properties of real asset funds: A survey of prior work and new findings. White paper, Institute for Private Capital, University of North Carolina at Chapel Hill, November 2025.

International Energy Agency. Energy and AI. World energy outlook special report, International Energy Agency, Paris, April 2025. URL `https://www.iea.org/reports/energy-and-ai`.

International Energy Agency. Key questions on energy and AI. Report, International Energy Agency, Paris, April 2026. URL `https://www.iea.org/reports/key-questions-on-energy-and-ai`.

Shmuel Kandel and Robert F. Stambaugh. On correlations and inferences about mean-variance efficiency. *Journal of Financial Economics*, 18(1):61–90, 1987.

Owen Kay, Robert Reaser, and Reid Taylor. Processing power: The effect of data centers on wholesale electricity markets. Working Paper 2606, Federal Reserve Bank of Dallas, March 2026.

Finn E. Kydland and Edward C. Prescott. Time to build and aggregate fluctuations. *Econometrica*, 50(6):1345–1370, 1982.

David C. Ling and Andy Naranjo. Returns and information transmission dynamics in public and private real estate markets. *Real Estate Economics*, 43(1):163–208, 2015.

Meta Platforms. Meta announces joint venture with funds managed by Blue Owl Capital to develop Hyperion data center. Press release, October 21, 2025.

Morgan Stanley Research. Bridging a $1.5tr data center financing gap. Technical report, Morgan Stanley, July 2025.

Graeme Newell, Hsu Wen Peng, and Anthony De Francesco. The performance of unlisted infrastructure in investment portfolios. *Journal of Property Research*, 28(1):59–74, 2011.

Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.

Jesse Noffsinger, Mark Patel, and Pankaj Sachdeva. The cost of compute: A $7 trillion race to scale data centers. Technical report, McKinsey & Company, April 2025.

Joseph L. Pagliari, Jr. Another take on real estate's role in mixed-asset portfolio allocations. *Real Estate Economics*, 45(1):75–132, 2017.

Joseph L. Pagliari, Jr., Kevin A. Scherer, and Richard T. Monopoli. Public versus private real estate equities: A more refined, long-term comparison. *Real Estate Economics*, 33(1):147–187, 2005.

L'uboš Pástor and Pietro Veronesi. Technological revolutions and stock prices. *American Economic Review*, 99(4):1451–1483, 2009.

Maximilian Rödel and Christoph Rothballer. Infrastructure as hedge against inflation—fact or fantasy? *Journal of Alternative Investments*, 15(1):110–123, 2012.

Richard Roll. A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory. *Journal of Financial Economics*, 4(2):129–176, 1977.

Richard Roll and Stephen A. Ross. On the cross-sectional relation between expected returns and betas. *Journal of Finance*, 49(1):101–121, 1994.

Hannah Rubinton and Bontu Ankit Patro. Tracking AI's contribution to GDP growth. Federal Reserve Bank of St. Louis, On the Economy, January 12, 2026. URL `https://www.stlouisfed.org/on-the-economy/2026/jan/tracking-ai-contribution-gdp-growth`.

Phurichai Rungcharoenkitkul. The AI investment race. BIS Working Paper 1367, Bank for International Settlements, July 2026.

Jay Shanken. Multivariate proxies and asset pricing relations: Living with the Roll critique. *Journal of Financial Economics*, 18(1):91–110, 1987.

Junying Shen, Ding Li, Grace Tiantian Qiu, Vishv Jeet, Michelle Yu Teng, and Ki Cheong Wong. Building a better portfolio: Balancing performance and liquidity. Technical report, PGIM Institutional Advisory & Solutions and GIC Economics & Investment Strategy, April 2020. URL

https://www.gic.com.sg/thinkspace/long-term-investing/building-a-better-portfolio-balancing-performance-and-liquidity/.

Robert F. Stambaugh. On the exclusion of assets from tests of the two-parameter model: A sensitivity analysis. *Journal of Financial Economics*, 10(3):237–268, 1982.

Dean Takahashi and Seth Alexander. Illiquid alternative asset fund modeling. *Journal of Portfolio Management*, 28(2):90–100, 2002.

Stijn Van Nieuwerburgh. Financing the AI buildout. Working paper, Columbia Business School, October 2026. URL https://ssrn.com/abstract=6444363. Draft of October 2, 2026, prepared for the Brookings Papers on Economic Activity. SSRN 6444363.

Jessica A. Wachter and Jonathan Wachter. What investment data implies about the AI transition. NBER Working Paper 35290, National Bureau of Economic Research, June 2026.

Appendix A. Real-asset portfolio definitions

Table A1 lists the GICS sub-industries from which each of the fifteen real-asset comparable portfolios is formed, together with constituent counts at year-end 2015, 2019, 2022, and 2025. Each portfolio is the value-weighted return on the listed firms in the sub-industries shown, after removal of any firm assigned to the demand complex, the buildout tier, or the mixed group of Table 1. The GICS structure effective after the close of March 17, 2023 created Data Center REITs (60108050) as a distinct sub-industry, carved from Specialized REITs, and moved the tower operators to Telecom Tower REITs (60108030). We apply each firm's latest available Compustat GICS classification retrospectively, supplemented by explicit assignments for the data center and tower/fiber REIT portfolios. The data center REIT portfolio includes Digital Realty, Equinix, and Iron Mountain throughout the sample; DuPont Fabros (January 2009 to September 2017), CoreSite (September 2010 to December 2021), CyrusOne (January 2013 to March 2022), QTS (October 2013 to August 2021), and Fermi (October 2025 onward). These dates reflect the available CRSP trading histories rather than separately imposed historical membership restrictions. CoreSite is identified by its CRSP PERMNO to avoid including unrelated firms that subsequently used the same ticker. Switch is not included.

TABLE A1: Definitions of the Real-Asset Comparable Portfolios

Portfolio	Channel	GICS sub-industries	2015	2019	2022	2025
IPP & Renewables	Power and fuel	Independent Power Producers & Energy Traders (55105010); Renewable Electricity (55105020)	25	21	22	20
Electric Utilities	Power and fuel	Electric Utilities (55101010)	40	37	36	35
Regulated Utilities ex Electric	Power and fuel	Gas Utilities (55102010); Multi-Utilities (55103010); Water Utilities (55104010)	56	48	45	44
Midstream Energy	Power and fuel	Oil & Gas Storage & Transportation (10102040)	106	81	67	66
Integrated Oil & Gas	Power and fuel	Integrated Oil & Gas (10102010)	22	18	14	14
Data Center REITs	Equipment and shell	Data Center REITs (60108050), supplemented by explicit ticker assignments	7	6	3	4
Heavy Electrical Equipment	Equipment and shell	Heavy Electrical Equipment (20104020)	8	8	12	13
Electrical Components & Equipment	Equipment and shell	Electrical Components & Equipment (20104010)	48	39	75	71
Construction & Engineering	Construction and materials	Construction & Engineering (20103010)	35	35	39	56
Building Products	Construction and materials	Building Products (20102010)	39	40	45	41
Steel	Construction and materials	Steel (15104050)	30	31	30	30
Copper	Construction and materials	Copper (15104025)	4	3	5	4
Aluminum	Construction and materials	Aluminum (15104010)	7	6	6	5
Tower/Fiber REITs	Connectivity	Telecom Tower REITs (60108030), supplemented by explicit assignments for CSAL and UNIT	4	4	4	4
Traditional REITs	Placebo	Selected equity and mortgage REIT sub-industries, including Mortgage REITs (40204010)	172	186	196	188

Each portfolio is the value-weighted return on the listed firms in the sub-industries shown, with firms assigned to the AI demand complex, the AI buildout tier, or the mixed group removed under the one-home rule. GICS classifications are based on the latest available Compustat classification for each firm and applied retrospectively, with explicit assignments supplementing the Data Center REITs and Tower/Fiber REITs portfolios. Constituent counts report the number of distinct securities (CRSP PERMNOs) on the final trading day of 2015, 2019, 2022, and 2025, respectively.

Appendix B. Synthetic control details

Table A2 reports the parallel-trends tests that underlie the inference in Section 5. For each portfolio and factor, the null hypothesis is that the differential factor beta of the portfolio relative to its synthetic control is the same in the 2010–2017 estimation period as in the 2018–2019 baseline; the annual test uses yearly differential betas, and the two-year test uses four two-year blocks. The buildout columns are the ones on which the results of Table 3 rest, and none of the eight exposed sectors rejects parallel pretrends in the buildout beta at the 10 percent level under either test.

TABLE A2: Parallel-Trends Tests for Synthetic Control Factor Betas

	Ex-AI Market		Buildout		Demand		Treasury	
Portfolio	**Annual**	**2-Year**	**Annual**	**2-Year**	**Annual**	**2-Year**	**Annual**	**2-Year**
Traditional REITs	< 0.001	< 0.001	0.393	0.924	0.274	0.730	0.018	0.037
Tower/Fiber REITs	0.016	0.403	0.798	0.841	< 0.001	0.814	0.270	0.330
Data Center REITs	0.160	0.027	0.438	0.672	0.209	0.924	0.167	0.015
Regulated Utilities ex Electric	0.001	< 0.001	0.264	0.876	0.441	0.182	0.031	0.044
IPP & Renewables	< 0.001	< 0.001	0.680	0.543	0.430	0.164	0.262	0.446
Midstream Energy	0.384	0.132	0.876	0.631	0.221	0.039	0.759	0.750
Electric Utilities	0.007	0.004	0.824	0.923	0.617	0.182	0.018	0.005
Electrical Components & Equipment	0.113	0.021	0.016	0.339	< 0.001	< 0.001	< 0.001	0.024
Heavy Electrical Equipment	0.171	0.675	0.178	0.431	0.429	0.836	0.502	0.256
Construction & Engineering	0.370	0.349	0.303	0.480	0.087	0.465	0.465	0.445
Copper	< 0.001	0.062	0.129	0.201	< 0.001	0.252	0.294	0.460
Aluminum	0.613	0.335	0.530	0.799	0.039	0.526	0.329	0.497
Steel	< 0.001	0.005	0.079	0.504	0.276	0.871	0.013	0.006
Integrated Oil & Gas	0.100	0.092	0.662	0.819	0.139	0.042	0.980	0.862
Building Products	0.641	0.518	0.950	0.905	0.009	0.002	0.371	0.259

This table reports p-values from tests of parallel pretrends in target-minus-synthetic factor betas over the pre-treatment period. Annual tests compare annual differential factor betas from 2010–2017 with the pooled 2018–2019 baseline; two-year tests use the 2010–2011, 2012–2013, 2014–2015, and 2016–2017 periods relative to the 2018–2019 baseline. The null hypothesis is that the pre-period differential factor betas are jointly equal to the 2018–2019 differential beta.

Table A3 reports the estimated donor weights used to construct each portfolio's synthetic control, together with out-of-sample validation diagnostics. The donor weights are estimated over 2010–2017, while validation performance is evaluated over the held-out 2018–2019 period. Validation RMSE measures the root-mean-squared difference between weekly target and synthetic-control returns, while the validation buildout beta gap measures the difference between their estimated buildout factor betas during 2018–2019. A gap close to zero indicates that the synthetic control closely matches the target portfolio's buildout exposure in the validation period. The es-

timated donor weights are generally concentrated in a small number of industries, although their composition varies considerably across portfolios.

Validation fit is relatively strong for seven of the eight sectors identified as exposed in the baseline analysis: their weekly return RMSEs range from 0.015 to 0.024, and their absolute buildout beta gaps are below 0.10. Heavy Electrical Equipment is the principal exception, with a validation RMSE of 0.050 and a buildout beta gap of 0.221. Several non-exposed portfolios exhibit larger buildout beta mismatches, particularly Copper and Steel. Overall, the validation results indicate that the synthetic controls provide reasonably close out-of-sample matches for most exposed sectors, although matching quality is not uniform. These diagnostics complement the parallel-trends tests by assessing the level of pre-treatment similarity rather than the stability of differential factor exposures over time.

TABLE A3: Synthetic-Control Donor Weights and Validation Fit

	Donor Weights (%)								Validation Fit	
Portfolio	GICS3010	GICS3020	GICS3030	GICS3510	GICS3520	GICS4010	GICS4020	GICS4030	RMSE	Buildout Gap
Heavy Electrical Equipment	13.4%	14.5%	0.0%	0.0%	0.0%	0.0%	72.1%	0.0%	0.050	0.221
IPP & Renewables	0.0%	54.6%	0.0%	0.0%	36.9%	0.0%	0.0%	8.5%	0.015	0.093
Construction & Engineering	0.0%	0.0%	0.0%	0.0%	0.0%	23.0%	57.9%	19.2%	0.016	0.063
Data Center REITs	0.0%	56.7%	0.0%	27.3%	15.9%	0.0%	0.0%	0.0%	0.024	-0.079
Electric Utilities	17.9%	15.6%	66.4%	0.0%	0.0%	0.0%	0.0%	0.0%	0.017	-0.078
Building Products	12.1%	0.0%	0.0%	0.0%	16.3%	5.9%	64.8%	0.9%	0.015	-0.002
Midstream Energy	41.1%	14.7%	0.0%	0.0%	21.9%	22.4%	0.0%	0.0%	0.018	-0.005
Regulated Utilities ex Electric	19.4%	43.9%	36.7%	0.0%	0.0%	0.0%	0.0%	0.0%	0.016	-0.079
Integrated Oil & Gas	25.8%	28.0%	0.0%	0.0%	0.0%	42.2%	0.0%	4.0%	0.020	0.024
Electrical Components & Equipment	20.3%	0.0%	0.0%	0.0%	7.8%	0.0%	53.1%	18.8%	0.016	0.235
Traditional REITs	13.4%	40.1%	0.0%	0.0%	20.7%	25.7%	0.0%	0.0%	0.017	-0.052
Copper	16.8%	0.0%	0.0%	0.0%	10.9%	45.1%	27.2%	0.0%	0.037	0.445
Aluminum	0.0%	0.0%	0.0%	0.0%	0.9%	50.7%	48.4%	0.0%	0.038	0.159
Tower/Fiber REITs	0.0%	86.1%	0.0%	6.8%	7.1%	0.0%	0.0%	0.0%	0.024	-0.026
Steel	0.0%	0.0%	0.0%	0.0%	0.0%	42.1%	57.9%	0.0%	0.028	0.273

This table reports the estimated donor weights used to construct the synthetic control for each portfolio, together with out-of-sample validation diagnostics. The first eight rows correspond to the sectors identified as exposed in the baseline analysis. Donor weights are estimated using 2010–2017 data and reported as percentages. The donor pool comprises eight non-AI GICS industry groups: Consumer Staples Distribution & Retail (3010), Food, Beverage & Tobacco (3020), Household & Personal Products (3030), Health Care Equipment & Services (3510), Pharmaceuticals, Biotechnology & Life Sciences (3520), Banks (4010), Financial Services (4020), and Insurance (4030). Validation statistics are calculated over the held-out 2018–2019 period. RMSE is the root-mean-squared difference between weekly target and synthetic-control returns, expressed in decimal return units. Buildout Gap is the target-minus-synthetic-control difference in estimated buildout factor betas.

Appendix C. Robustness of the difference-in-differences estimates

Tables A4, A5, and A6 report the corresponding difference-in-differences estimates for the ex-NVIDIA, equal-weighted, and common-control robustness specifications. The main pattern of increased buildout exposure persists, although statistical significance varies across specifications. Among the eight sectors identified as exposed in the baseline, seven retain FDR-adjusted significance under the ex-NVIDIA specification, four under equal weighting, and seven under the common control. Combining statistical significance with the corresponding parallel-trends tests leaves six, four, and four sectors, respectively, satisfying both criteria. Thus, the direction of the results is robust, although the strength of sector-level inference depends on factor construction and control selection.

TABLE A4: Changes in the Factor Exposures of the Public Comparables: Ex-NVIDIA Buildout Factor

Portfolio	**Buildout $\Delta\beta$**	**Demand $\Delta\beta$**	**Ex-AI Market $\Delta\beta$**	**Treasury $\Delta\beta$**
Heavy Electrical Equipment	$1.189^{***\dagger}$	0.180	-1.279	0.966
	(0.416)	(0.595)	(0.864)	(1.958)
IPP & Renewables	$0.689^{***\dagger}$	-0.081	-0.288	0.271
	(0.186)	(0.191)	(0.363)	(0.542)
Construction & Engineering	$0.704^{***\dagger}$	-0.427	-0.779^{**}	0.593
	(0.183)	(0.248)	(0.353)	(0.545)
Data Center REITs	$0.563^{***\dagger}$	-0.085	-0.487	-0.613
	(0.208)	(0.226)	(0.363)	(0.834)
Electric Utilities	$0.359^{***\dagger}$	-0.202	-0.120	0.178
	(0.115)	(0.185)	(0.303)	(0.435)
Building Products	$0.420^{**\dagger}$	-0.129	-0.428	-0.062
	(0.164)	(0.295)	(0.349)	(0.426)
Midstream Energy	0.212	-0.302	-0.253	-0.360
	(0.170)	(0.219)	(0.455)	(0.605)
Regulated Utilities ex Electric	$0.245^{**\dagger}$	-0.213	-0.018	0.044
	(0.104)	(0.170)	(0.306)	(0.357)
Integrated Oil & Gas	0.095	-0.493^{**}	-0.111	-0.669
	(0.167)	(0.214)	(0.438)	(0.765)
Electrical Components & Equipment	0.175	0.009	-0.249	0.499
	(0.120)	(0.129)	(0.220)	(0.482)
Traditional REITs	0.001	-0.119	0.428^{*}	-0.450
	(0.112)	(0.165)	(0.246)	(0.381)
Copper	0.043	-0.147	0.404	0.616
	(0.263)	(0.391)	(0.549)	(1.281)
Aluminum	0.104	0.026	-0.030	0.930
	(0.268)	(0.358)	(0.468)	(1.097)
Tower/Fiber REITs	0.050	-0.148	-0.028	0.029
	(0.223)	(0.310)	(0.562)	(0.951)
Steel	-0.073	-0.266	0.336	0.532
	(0.148)	(0.210)	(0.459)	(0.713)

This table reports the exposure change, $\Delta\beta$, for each public comparable portfolio and each factor, estimated as a synthetic control difference-in-differences in buildout, demand, ex-AI market, and Treasury factor betas. The buildout factor is value weighted and excludes NVIDIA. Synthetic controls are estimated using weekly data from 2010–2017, with 2018–2019 used as a held-out validation period and as the pre-period baseline. We omit 2020–2022 because the COVID-19 pandemic and the associated policy and market responses created unusually large disruptions that are difficult to disentangle from changes in underlying factor exposures. The exposure changes compare 2018–2019 with the post-period, January 2025–March 2026; observations from 2020–2024 are omitted from the pre/post comparison. Bootstrap standard errors are reported in parentheses. Inference is based on repeated resampling of blocks of weekly returns, with the synthetic control rebuilt in each draw to account for uncertainty in both returns and the construction of the synthetic counterfactual. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$ based on unadjusted bootstrap p-values. † denotes estimates that remain significant at the 10% level after Benjamini–Hochberg false-discovery-rate adjustment across portfolios. Among the eight portfolios shown above the line, six retain buildout exposure changes that remain significant at the 10% level after false-discovery-rate adjustment and also pass both annual and two-year buildout pretrend tests.

TABLE A5: Changes in the Factor Exposures of the Public Comparables: Equal-Weighted Buildout Factor

Portfolio	**Buildout** $\Delta\beta$	**Demand** $\Delta\beta$	**Ex-AI Market** $\Delta\beta$	**Treasury** $\Delta\beta$
Heavy Electrical Equipment	0.692	0.583	−0.939	0.696
	(0.483)	(0.570)	(0.890)	(1.634)
IPP & Renewables	$0.679^{***\dagger}$	0.030	−0.521	0.191
	(0.151)	(0.169)	(0.345)	(0.523)
Construction & Engineering	$0.479^{**\dagger}$	−0.184	-0.780^{**}	0.434
	(0.167)	(0.224)	(0.360)	(0.466)
Data Center REITs	$0.545^{***\dagger}$	−0.042	−0.576	−0.594
	(0.205)	(0.195)	(0.394)	(0.776)
Electric Utilities	$0.342^{***\dagger}$	−0.129	−0.202	0.151
	(0.118)	(0.189)	(0.288)	(0.432)
Building Products	0.288^{*}	0.015	−0.325	−0.116
	(0.148)	(0.274)	(0.346)	(0.401)
Midstream Energy	0.164	−0.266	−0.273	−0.342
	(0.152)	(0.217)	(0.442)	(0.573)
Regulated Utilities ex Electric	0.216^{**}	−0.143	−0.044	−0.010
	(0.108)	(0.166)	(0.285)	(0.364)
Integrated Oil & Gas	0.164	-0.514^{**}	−0.159	−0.546
	(0.140)	(0.207)	(0.404)	(0.738)
Electrical Components & Equipment	0.058	0.134	−0.166	0.358
	(0.116)	(0.110)	(0.221)	(0.390)
Traditional REITs	−0.047	−0.105	0.482^{*}	−0.479
	(0.099)	(0.165)	(0.250)	(0.370)
Copper	−0.278	0.107	0.615	0.401
	(0.224)	(0.346)	(0.543)	(1.316)
Aluminum	−0.205	0.134	0.286	0.743
	(0.217)	(0.273)	(0.522)	(1.054)
Tower/Fiber REITs	−0.038	−0.096	0.057	−0.052
	(0.181)	(0.295)	(0.501)	(0.866)
Steel	−0.200	−0.166	0.396	0.297
	(0.143)	(0.222)	(0.491)	(0.716)

This table reports the exposure change, $\Delta\beta$, for each public comparable portfolio and each factor, estimated as a synthetic control difference-in-differences in buildout, demand, ex-AI market, and Treasury factor betas. The buildout factor is equal weighted across its constituent firms. Synthetic controls are estimated using weekly data from 2010–2017, with 2018–2019 used as a held-out validation period and as the pre-period baseline. We omit 2020–2022 because the COVID-19 pandemic and the associated policy and market responses created unusually large disruptions that are difficult to disentangle from changes in underlying factor exposures. The exposure changes compare 2018–2019 with the post-period, January 2025–March 2026; observations from 2020–2024 are omitted from the pre/post comparison. Bootstrap standard errors are reported in parentheses. Inference is based on repeated resampling of blocks of weekly returns, with the synthetic control rebuilt in each draw to account for uncertainty in both returns and the construction of the synthetic counterfactual. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$ based on unadjusted bootstrap p-values. † denotes estimates that remain significant at the 10% level after Benjamini–Hochberg false-discovery-rate adjustment across portfolios. Among the eight portfolios shown above the line, four retain buildout exposure changes that remain significant at the 10% level after false-discovery-rate adjustment and also pass both annual and two-year buildout pretrend tests.

TABLE A6: Changes in the Factor Exposures of the Public Comparables: Common Control

Portfolio	**Buildout** $\Delta\beta$	**Demand** $\Delta\beta$	**Ex-AI Market** $\Delta\beta$	**Treasury** $\Delta\beta$
Heavy Electrical Equipment	0.690^{*}	0.509	-0.602	1.107
	(0.358)	(0.594)	(0.805)	(1.785)
IPP & Renewables	$0.731^{***\dagger}$	-0.291^{*}	-0.215	0.500
	(0.130)	(0.151)	(0.340)	(0.632)
Construction & Engineering	$0.469^{**\dagger}$	-0.193	-0.356	0.966^{**}
	(0.163)	(0.275)	(0.324)	(0.442)
Data Center REITs	$0.528^{***\dagger}$	-0.239	-0.445	-0.683
	(0.182)	(0.234)	(0.289)	(0.845)
Electric Utilities	$0.409^{***\dagger}$	$-0.296^{**\dagger}$	-0.153	0.055
	(0.067)	(0.101)	(0.243)	(0.621)
Building Products	$0.282^{**\dagger}$	-0.027	-0.162	0.116
	(0.130)	(0.290)	(0.343)	(0.379)
Midstream Energy	$0.348^{***\dagger}$	$-0.494^{***\dagger}$	-0.213	-0.086
	(0.096)	(0.151)	(0.398)	(0.481)
Regulated Utilities ex Electric	$0.320^{***\dagger}$	$-0.376^{***\dagger}$	-0.072	-0.113
	(0.074)	(0.097)	(0.229)	(0.540)
Integrated Oil & Gas	0.187	$-0.588^{***\dagger}$	-0.041	-0.404
	(0.122)	(0.169)	(0.387)	(0.681)
Electrical Components & Equipment	0.077	0.139	-0.094	0.569
	(0.100)	(0.134)	(0.175)	(0.393)
Traditional REITs	0.116	$-0.265^{**\dagger}$	0.371^{*}	-0.399
	(0.077)	(0.106)	(0.191)	(0.342)
Copper	-0.018	0.062	0.648	0.999
	(0.201)	(0.316)	(0.476)	(1.286)
Aluminum	0.003	0.218	0.262	1.343
	(0.230)	(0.327)	(0.426)	(1.127)
Tower/Fiber REITs	0.131	-0.346	-0.162	-0.229
	(0.185)	(0.260)	(0.456)	(1.000)
Steel	-0.182	-0.061	0.603	0.827
	(0.125)	(0.206)	(0.421)	(0.678)

This table reports the exposure change, $\Delta\beta$, for each public comparable portfolio and each factor, estimated using a common-control difference-in-differences specification in buildout, demand, ex-AI market, and Treasury factor betas. Rather than estimating portfolio-specific synthetic controls, all portfolios are compared against the same value-weighted control portfolio comprising non-AI firms in Consumer Staples, Health Care, and Financials. The analysis uses weekly returns, with 2018–2019 serving as the pre-period baseline. We omit 2020–2022 because the COVID-19 pandemic and the associated policy and market responses created unusually large disruptions that are difficult to disentangle from changes in underlying factor exposures. The exposure changes compare 2018–2019 with the post-period, January 2025–March 2026; observations from 2020–2024 are omitted from the pre/post comparison. Bootstrap standard errors are reported in parentheses. Inference is based on repeated resampling of blocks of weekly returns, without estimating or rebuilding synthetic control weights in each draw. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$ based on unadjusted bootstrap p-values. † denotes estimates that remain significant at the 10% level after Benjamini–Hochberg false-discovery-rate adjustment across portfolios. Among the eight portfolios shown above the line, seven retain buildout exposure changes that remain significant at the 10% level after false-discovery-rate adjustment, of which four also pass both annual and two-year buildout pretrend tests.

Tables A7, A8, and A9 report the corresponding parallel-trends tests for the ex-NVIDIA, equal-weighted, and common-control robustness specifications. Under the ex-NVIDIA factor, seven of the eight exposed sectors satisfy both tests, with Regulated Utilities ex Electric rejecting parallel pretrends in the annual test. Under equal weighting, all eight sectors satisfy both tests. The common-control specification produces somewhat weaker evidence: Regulated Utilities ex Electric, Construction & Engineering, and Electric Utilities reject parallel pretrends in the annual test, although none rejects in the two-year test. Thus, seven of eight sectors satisfy both tests under ex-NVIDIA, eight of eight under equal weighting, and five of eight under the common-control specification.

TABLE A7: Parallel-Trends Tests for Synthetic Control Factor Betas, Ex-NVIDIA

	Ex-AI Market		**Buildout**		**Demand**		**Treasury**	
Portfolio	**Annual**	**2-Year**	**Annual**	**2-Year**	**Annual**	**2-Year**	**Annual**	**2-Year**
Traditional REITs	< 0.001	< 0.001	0.445	0.899	0.326	0.757	0.018	0.039
Tower/Fiber REITs	0.019	0.446	0.878	0.807	< 0.001	0.806	0.280	0.310
Data Center REITs	0.182	0.031	0.436	0.604	0.206	0.903	0.178	0.018
Regulated Utilities ex Electric	0.003	0.004	0.072	0.821	0.438	0.254	0.042	0.055
IPP & Renewables	< 0.001	< 0.001	0.792	0.634	0.305	0.133	0.288	0.502
Midstream Energy	0.470	0.145	0.893	0.735	0.229	0.055	0.791	0.721
Electric Utilities	0.026	0.024	0.804	0.939	0.652	0.220	0.020	0.006
Electrical Components & Equipment	0.079	0.026	0.003	0.254	< 0.001	< 0.001	< 0.001	0.012
Heavy Electrical Equipment	0.159	0.665	0.132	0.408	0.390	0.832	0.465	0.260
Construction & Engineering	0.268	0.392	0.253	0.460	0.107	0.409	0.514	0.503
Copper	< 0.001	0.041	0.165	0.144	< 0.001	0.313	0.287	0.498
Aluminum	0.530	0.407	0.595	0.755	0.027	0.531	0.399	0.502
Steel	< 0.001	0.008	0.644	0.604	0.247	0.826	0.006	0.004
Integrated Oil & Gas	0.112	0.111	0.923	0.812	0.143	0.064	0.979	0.866
Building Products	0.627	0.465	0.935	0.652	0.026	0.007	0.549	0.330

This table reports p-values from tests of parallel pretrends in target-minus-synthetic factor betas over the pre-treatment period, using a value-weighted AI buildout factor that excludes NVIDIA. Annual tests compare annual differential factor betas from 2010–2017 with the pooled 2018–2019 baseline; two-year tests use the 2010–2011, 2012–2013, 2014–2015, and 2016–2017 periods relative to the 2018–2019 baseline. The null hypothesis is that the pre-period differential factor betas are jointly equal to the 2018–2019 differential beta.

TABLE A8: Parallel-Trends Tests for Synthetic Control Factor Betas, Equal-Weight

	Ex-AI Market		Buildout		Demand		Treasury	
Portfolio	**Annual**	**2-Year**	**Annual**	**2-Year**	**Annual**	**2-Year**	**Annual**	**2-Year**
Traditional REITs	0.002	< 0.001	0.411	0.991	0.329	0.703	0.025	0.050
Tower/Fiber REITs	0.071	0.548	0.895	0.999	< 0.001	0.864	0.270	0.338
Data Center REITs	0.881	0.232	0.636	0.914	0.121	0.751	0.189	0.021
Regulated Utilities ex Electric	0.007	0.007	0.321	0.981	0.665	0.431	0.051	0.067
IPP & Renewables	< 0.001	< 0.001	0.707	0.448	0.260	0.109	0.211	0.342
Midstream Energy	0.203	0.086	0.390	0.123	0.123	0.025	0.767	0.564
Electric Utilities	0.048	0.029	0.471	0.876	0.696	0.240	0.007	0.004
Electrical Components & Equipment	0.052	0.004	< 0.001	0.009	< 0.001	< 0.001	< 0.001	< 0.001
Heavy Electrical Equipment	0.243	0.696	0.562	0.723	0.354	0.749	0.329	0.157
Construction & Engineering	0.291	0.293	0.265	0.403	0.067	0.410	0.341	0.303
Copper	< 0.001	0.004	0.104	0.014	< 0.001	0.289	0.427	0.515
Aluminum	0.507	0.293	0.428	0.877	0.003	0.645	0.184	0.547
Steel	0.004	0.023	0.115	0.693	0.306	0.949	0.039	0.007
Integrated Oil & Gas	0.147	0.140	0.786	0.780	0.152	0.064	0.979	0.885
Building Products	0.795	0.256	0.881	0.861	0.006	0.001	0.419	0.254

This table reports p-values from tests of parallel pretrends in target-minus-synthetic factor betas over the pre-treatment period, using an equal-weighted AI buildout factor. Annual tests compare annual differential factor betas from 2010–2017 with the pooled 2018–2019 baseline; two-year tests use the 2010–2011, 2012–2013, 2014–2015, and 2016–2017 periods relative to the 2018–2019 baseline. The null hypothesis is that the pre-period differential factor betas are jointly equal to the 2018–2019 differential beta.

TABLE A9: Parallel-Trends Tests for Factor Betas, Common Control

	Ex-AI Market		Buildout		Demand		Treasury	
Portfolio	**Annual**	**2-Year**	**Annual**	**2-Year**	**Annual**	**2-Year**	**Annual**	**2-Year**
Traditional REITs	0.005	0.003	0.202	0.787	0.339	0.953	0.004	0.012
Tower/Fiber REITs	0.097	0.784	0.768	0.594	< 0.001	0.914	0.081	0.051
Data Center REITs	0.541	0.068	0.282	0.548	0.481	0.782	0.019	0.001
Regulated Utilities ex Electric	0.003	0.011	0.055	0.104	0.987	0.983	< 0.001	< 0.001
IPP & Renewables	< 0.001	< 0.001	0.529	0.467	0.295	0.049	0.011	0.010
Midstream Energy	0.377	0.108	0.878	0.527	0.450	0.138	0.872	0.770
Electric Utilities	0.002	0.022	0.048	0.144	0.962	0.817	< 0.001	< 0.001
Electrical Components & Equipment	0.316	0.033	0.003	0.094	< 0.001	< 0.001	0.010	0.010
Heavy Electrical Equipment	0.136	0.592	0.170	0.410	0.425	0.847	0.370	0.149
Construction & Engineering	0.056	0.258	0.015	0.141	0.066	0.831	0.196	0.078
Copper	< 0.001	0.022	0.097	0.207	0.003	0.330	0.244	0.242
Aluminum	0.430	0.174	0.783	0.942	0.015	0.696	0.256	0.204
Steel	< 0.001	< 0.001	0.108	0.368	0.038	0.824	0.018	0.005
Integrated Oil & Gas	0.115	0.094	0.605	0.834	0.300	0.186	0.978	0.779
Building Products	0.706	0.531	0.963	0.969	0.044	0.012	0.866	0.411

This table reports p-values from tests of parallel pretrends in target-minus-control factor betas over the pre-treatment period, using a common value-weighted control portfolio comprising non-AI firms in consumer staples, health care, and financials, rather than portfolio-specific synthetic controls. Annual tests compare annual differential factor betas from 2010–2017 with the pooled 2018–2019 baseline; two-year tests use the 2010–2011, 2012–2013, 2014–2015, and 2016–2017 periods relative to the 2018–2019 baseline. The null hypothesis is that the pre-period differential factor betas are jointly equal to the 2018–2019 differential beta.

Table A10 extends the donor-weight and validation-fit analysis to the three robustness specifications. Excluding NVIDIA from the buildout factor produces donor allocations broadly similar to those in the baseline specification, while equal weighting generates more substantial changes in the composition of several synthetic controls. The common-control specification instead uses a single value-weighted benchmark, so portfolio-specific donor weights are not estimated. Despite these differences, validation RMSEs remain relatively stable across specifications. Seven of the eight exposed sectors have validation RMSEs below 0.030 in all three specifications, while Heavy Electrical Equipment consistently exhibits substantially higher tracking error. The validation buildout beta gaps display greater sensitivity to the specification, particularly under equal weighting, where the gap for Heavy Electrical Equipment rises to 0.625. The common-control specification also produces larger beta mismatches for some sectors, although the gaps remain relatively modest for most. Overall, the robustness specifications preserve broadly similar out-of-sample return fit, but the precision with which controls reproduce pre-treatment buildout factor exposures is more sensitive to factor construction and control selection.

TABLE A10: Synthetic-Control Donor Weights and Validation Fit: Robustness Specifications

	Donor Weights (%)								Validation Fit	
Panel A: Ex-NVIDIA Buildout Factor	GICS3010	GICS3020	GICS3030	GICS3510	GICS3520	GICS4010	GICS4020	GICS4030	RMSE	Buildout Gap
Heavy Electrical Equipment	0.0%	25.8%	0.0%	0.0%	0.0%	0.0%	74.2%	0.0%	0.050	0.116
IPP & Renewables	0.0%	52.3%	0.0%	0.0%	44.9%	0.0%	0.0%	2.8%	0.016	0.075
Construction & Engineering	0.0%	0.0%	0.0%	0.0%	0.0%	30.2%	56.2%	13.6%	0.016	0.027
Data Center REITs	0.0%	58.8%	0.0%	28.6%	12.6%	0.0%	0.0%	0.0%	0.024	-0.104
Electric Utilities	18.4%	16.5%	65.1%	0.0%	0.0%	0.0%	0.0%	0.0%	0.017	-0.095
Building Products	9.9%	0.0%	6.5%	0.0%	13.7%	13.8%	56.0%	0.0%	0.014	-0.039
Midstream Energy	30.0%	23.2%	0.0%	0.0%	20.3%	26.5%	0.0%	0.0%	0.018	-0.032
Regulated Utilities ex Electric	20.1%	47.7%	32.3%	0.0%	0.0%	0.0%	0.0%	0.0%	0.016	-0.082
Integrated Oil & Gas	19.7%	30.3%	0.0%	0.0%	0.0%	41.5%	0.0%	8.4%	0.020	0.005
Electrical Components & Equipment	18.4%	0.0%	0.0%	0.0%	10.0%	0.0%	59.6%	12.0%	0.016	0.231
Traditional REITs	16.3%	37.7%	0.0%	0.0%	21.2%	24.5%	0.0%	0.3%	0.017	-0.022
Copper	21.5%	0.0%	0.0%	0.0%	9.0%	64.2%	5.3%	0.0%	0.037	0.505
Aluminum	0.0%	0.0%	0.0%	0.0%	7.0%	65.4%	27.6%	0.0%	0.038	0.166
Tower/Fiber REITs	0.0%	85.6%	0.0%	10.3%	4.0%	0.0%	0.0%	0.0%	0.024	-0.002
Steel	2.5%	0.0%	0.0%	0.0%	0.0%	48.6%	49.0%	0.0%	0.028	0.253
Panel B: Equal-Weighted Buildout Factor	GICS3010	GICS3020	GICS3030	GICS3510	GICS3520	GICS4010	GICS4020	GICS4030	RMSE	Buildout Gap
Heavy Electrical Equipment	0.0%	0.0%	0.0%	0.0%	59.6%	0.0%	10.2%	30.2%	0.050	0.625
IPP & Renewables	0.0%	52.7%	0.0%	0.0%	28.5%	0.0%	0.0%	18.7%	0.015	0.146
Construction & Engineering	0.0%	0.0%	0.0%	0.0%	0.0%	20.8%	79.2%	0.0%	0.016	0.102
Data Center REITs	0.0%	44.9%	0.0%	55.1%	0.0%	0.0%	0.0%	0.0%	0.025	-0.031
Electric Utilities	18.4%	17.9%	61.8%	0.0%	1.9%	0.0%	0.0%	0.0%	0.017	-0.088
Building Products	0.0%	0.0%	0.0%	0.0%	24.6%	0.0%	54.4%	21.0%	0.015	0.045
Midstream Energy	38.1%	32.7%	0.0%	0.0%	2.1%	27.0%	0.0%	0.0%	0.019	0.022
Regulated Utilities ex Electric	15.3%	48.5%	24.8%	0.0%	11.4%	0.0%	0.0%	0.0%	0.016	-0.078
Integrated Oil & Gas	30.9%	17.2%	0.0%	0.0%	0.0%	23.1%	0.0%	28.9%	0.020	0.028
Electrical Components & Equipment	0.0%	0.0%	0.0%	0.0%	28.5%	0.0%	44.5%	27.0%	0.016	0.335
Traditional REITs	18.0%	38.0%	0.0%	0.0%	17.8%	26.1%	0.0%	0.0%	0.017	-0.015
Copper	43.1%	0.0%	0.0%	0.0%	0.0%	56.9%	0.0%	0.0%	0.038	0.650
Aluminum	0.0%	0.0%	0.0%	0.0%	17.7%	79.3%	3.1%	0.0%	0.038	0.437
Tower/Fiber REITs	0.0%	78.3%	0.0%	8.9%	12.8%	0.0%	0.0%	0.0%	0.024	0.011
Steel	0.0%	0.0%	0.0%	0.0%	0.0%	64.5%	35.5%	0.0%	0.028	0.262
Panel C: Common Control	GICS3010	GICS3020	GICS3030	GICS3510	GICS3520	GICS4010	GICS4020	GICS4030	RMSE	Buildout Gap
Heavy Electrical Equipment									0.049	0.251
IPP & Renewables									0.016	0.040
Construction & Engineering									0.017	0.100
Data Center REITs									0.027	-0.123
Electric Utilities									0.022	-0.094
Building Products									0.014	0.030
Midstream Energy									0.018	-0.049
Regulated Utilities ex Electric									0.021	-0.119
Integrated Oil & Gas									0.018	-0.007
Electrical Components & Equipment									0.017	0.245
Traditional REITs									0.018	-0.090
Copper									0.039	0.458
Aluminum									0.040	0.205
Tower/Fiber REITs									0.027	-0.088
Steel									0.031	0.324

This table reports donor weights and out-of-sample validation diagnostics for three robustness specifications. Panel A excludes NVIDIA from the buildout factor, Panel B uses an equal-weighted buildout factor, and Panel C replaces portfolio-specific synthetic controls with a common value-weighted control portfolio comprising non-AI firms in consumer staples, health care, and financials. The first eight rows of each panel are the sectors identified as exposed in the baseline analysis. Donor weights are estimated over 2010–2017 and reported as percentages; the eight GICS donor industries are defined in Table A3. Donor weights are not estimated in Panel C and are therefore left blank. Validation statistics are calculated over 2018–2019. RMSE measures the root-mean-squared difference between weekly target and control returns in decimal units. Buildout Gap measures the target-minus-control difference in estimated buildout factor betas.

Appendix D. Derivation for Section 4

Write the measured beta of equation (4) as $\beta_s = N/D$, with $N = \beta_s^O b_O \sigma_O^2 + \beta_s^A b_A \sigma_A^2$ and $D = b_O^2 \sigma_O^2 + b_A^2 \sigma_A^2$; differentiating with respect to b_A gives $\partial \beta_s / \partial b_A = \sigma_A^2 \left(\beta_s^A D - 2 b_A N \right) / D^2$. The derivative is positive if and only if $\beta_s^A > 2 b_A \, N/D = 2 b_A \beta_s$, which is the threshold in equation (5).

Appendix E. Fund reconstruction

Deal counts are known for some funds but not others. For some of the funds with a known number of deals, not all deals are accounted for, and these partially observed, or "dirty", funds constitute the majority of the sample.[22] Accordingly, a comprehensive analysis cannot rely entirely on "clean" funds, that is, those with a known number of deals and for which all are accounted for in the data. Compared with dirty funds, clean funds are not a random subset of the sample: they contain smaller deals, have a lower infrastructure share and lower buildout exposure, and differ substantially in both sector composition and, within real estate, industry composition. At the fund level, dirty funds are also larger by observed invested capital, reported deal count, and fund size, and are somewhat older on average.

To approach a comprehensive and representative fund-level analysis, we fill in missing deals in dirty funds with deals from funds that qualify as neither clean nor dirty, specifically, those for which a reported investment count is unavailable, those with fewer than three reported investments or fewer than three observed deals, and those whose observed investments span more than seven years. We call these "fundless" deals.

The resulting sample contains 159 clean funds with 2,680 observed deals and 530 dirty funds with 7,478 observed deals, together with 2,020 fundless deals. The dirty funds are missing 8,285 deals relative to their reported investment counts, substantially more than can be supplied by the fundless pool alone. Because there are not enough fundless deals to complete all dirty funds, the reconstruction also releases the observed deals of dirty funds that are not retained in a given reconstruction and makes those deals available for reassignment to other dirty funds.

The reconstruction procedure uses clean funds to characterize fund membership and then ap-

[22] When the number of observed deals exceeds the reported investment count, we treat the observed count as the fund's deal count and classify the fund as clean. This holds for 11 funds with vintages before 2010, 4 funds with vintages between 2010 and 2019, and 7 funds with vintages since 2020.

plies those relationships to complete dirty funds. The procedure is deliberately conservative: candidate deals must satisfy the same timing rules used in model estimation, reconstructed funds are screened against the composition of observed complete funds, and downstream analysis is summarized across multiple high-quality reconstructions rather than relying on a single completed panel.

Model estimation uses artificial missing deal exercises constructed from clean funds. One clean fund deal is removed, the remaining deals define the observed fund, and the removed deal is appended to a candidate set drawn from the same pool used in reconstruction: observed dirty fund deals plus fundless deals. This process is repeated in turn for every deal in every clean fund, so that each deal serves once as the held-out true member while the other deals in its fund provide the observed fund context.[23] Candidate eligibility is anchored to the target fund's earliest observed entry date (i.e., our definition of the fund's vintage), with both the held-out deal and competing candidates required to enter on or after that date and no more than seven years later.

Each candidate is described by six fit measures: entry date, GICS sector, GICS subsector, real asset subsector, investment size, and gross TVM. For each of the three categorical characteristics, the candidate's fit score is the fraction of deals already in the fund that belong to the same category. A candidate in a category that is common within the fund therefore receives a high score, while a candidate in a category not represented in the fund receives a score of zero. For entry date, investment size, and gross TVM, the procedure first measures how far each existing fund deal is from its closest other fund deal on that characteristic. It then measures the candidate's distance from its closest existing fund deal and scores the candidate according to how that distance compares with the fund's own within-fund distances. A candidate therefore receives a high score when it is at least as close to the fund as the fund's existing deals typically are to one another. Entry date distances are measured in days, investment size is compared on a logarithmic scale, and gross TVM is compared in its original scale.

The six measures enter a conditional multinomial logit model with nonnegative coefficients. Each artificial missing deal exercise is a separate choice problem in which the held-out clean fund deal is the observed choice and the candidate pool deals are alternatives. The estimated coefficients determine how strongly each dimension of similarity helps distinguish actual fund mem-

[23] Clean fund deals are not competing alternatives because their memberships are treated as fixed.

bers. We evaluate the model out of sample using five-fold cross-validation. Clean funds are divided into five folds, and the model is estimated five times, each time using 80 percent of the funds for estimation and the remaining 20 percent for evaluation. Each fund therefore appears in the held-out sample once, and performance is summarized across all held-out observations. Out of sample, the true held-out deal is ranked first about 12 percent of the time, compared with only about 0.06 percent under random assignment. Across experiments, there are roughly 2,700 eligible candidates on average, after restricting the candidate set to deals available to the reconstruction procedure that enter between the target fund's earliest observed investment date and seven years afterward. The random benchmark accounts for variation in the number of eligible candidates across experiments. The true deal is also ranked in the top 10 candidates about 28 percent of the time, compared with about 0.6 percent under random assignment, and has a median rank of 50, placing it within roughly the top 2.1 percent of the eligible candidate set. The model selects the true deal at roughly 198 times the random top 1 rate and 47 times the random top 10 rate. These results indicate that the model has substantial discriminatory power: it identifies the true fund member far more often than chance and generally ranks it near the top of a candidate set containing thousands of plausible alternatives.

Clean funds are then used to determine how well the deals in an actual fund fit together according to the matching model. For each clean fund, one deal is removed at a time and scored against the eligible candidate pool based on how well it fits the remaining deals in that fund. This score is converted into a membership percentile, defined as the share of eligible outside candidates receiving a lower matching score, so that higher percentiles indicate stronger fit. For each fund, we summarize these membership percentiles using the 25th percentile, so the test reflects how well the weaker-fitting deals belong rather than being driven by a few particularly strong matches. We then examine this measure across all clean funds and use its fifth percentile as the cutoff. A reconstructed fund is rejected as dubious if its lower-quartile membership percentile falls below this benchmark, that is, if its weaker-fitting deals fit together worse than those of all but the bottom 5 percent of observed clean funds.

We generate 1,000 reconstruction attempts using different random sequences. Within each attempt, a feasible subset of dirty funds is selected subject to the stock of deals available for reconstruction. For each retained fund, missing deals are filled one deal at a time, with the order

of incomplete funds randomized before each assignment. When a fund is selected, its highest scoring eligible candidate is assigned to it and removed from the common pool. After each assignment, the deal becomes part of the fund and the remaining candidates are rescored using the expanded fund. If the selected set of dirty funds cannot be completed under the inventory and timing constraints, problematic funds can be dropped rather than forcing an implausible allocation (for example, a deal with an entry date 15 years after the fund's first entry date).

Each completed dirty fund is then evaluated using the clean fund plausibility benchmark described above, and funds that fall below the cutoff are removed. Among successful reconstruction attempts, the reconstruction attempts are ranked after this screening, first by the number of completed dirty funds that remain and then by the mean matching score of their surviving assignments. The top twenty screened reconstructions are retained for fund-level analysis. Each panel contains the original clean funds together with reconstructed dirty funds that passed the screen, and subsequent analyses are run separately in all twenty panels and summarized across them. The procedure does not require every dirty fund or candidate deal to be resolved when the available information does not support a sufficiently strong assignment.

Appendix F. Synthetic NAV construction

The deal-level data contain investment entry dates, exit or report dates, invested capital, and terminal value information, but they do not provide a complete time series of NAV for each investment. Several analyses nevertheless require a measure of how much capital is economically exposed during each calendar year. We therefore construct a synthetic NAV path for every investment over its observed life. The objective is to create a consistent exposure measure that begins at invested capital, evolves smoothly over the holding period, and ends at the amount of value that remains unrealized, rather than to infer a market price at every point in time.

For each investment, the synthetic NAV path runs from its entry date to its exit date. For investments that have not exited, the report date is used as the terminal date. At the terminal date, gross value is measured as realized plus unrealized value when both are available. If not, we use total value, followed by invested capital multiplied by gross TVM, and finally invested capital as a fallback. Remaining unrealized value is taken directly when reported; otherwise, if realized value is known, it is calculated as gross value less realized value.

The procedure then constructs two smooth value paths over the investment's observed life. Let f denote the fraction of the holding period that has elapsed, ranging from zero at entry to one at the terminal date. The first path is a gross value path, $G(f)$, which begins at invested capital and grows geometrically to gross terminal value. The second is an unrealized value path, $U(f)$, which also begins at invested capital but grows geometrically to terminal unrealized value. When the relevant terminal value is zero, the corresponding path declines linearly from invested capital to zero instead. This treatment avoids taking powers of zero while preserving the intended endpoint.

Synthetic NAV is a weighted average of these two paths. The weight on the gross value path is given by $w(f) := 1 - f^k$, and synthetic NAV is given by

$$\text{NAV}(f) := w(f)G(f) + [1 - w(f)]\, U(f).$$

The exponent k controls the shape of the transition. With $k = 1$, the shift from gross value to unrealized value is linear over the investment life. Values above one keep more weight on gross value during the earlier portion of the holding period and move more of the transition toward the later portion. Our baseline $k = 3$ therefore produces a moderately bowed, back-loaded realization pattern: the investment remains closer to its gross value path early in life, while realized value increasingly reduces NAV as the terminal date approaches. The parameter is a modeling choice rather than an estimated structural parameter, and the analysis can be recomputed under alternative values of k as a sensitivity check.

At entry, $f = 0$ and the weight on the gross path is one, so synthetic NAV equals invested capital. As the investment ages, weight gradually shifts from the gross value path toward the unrealized value path. At the terminal date, $f = 1$ and the gross path weight is zero, so synthetic NAV equals terminal unrealized value. A fully realized investment therefore ends at zero NAV, while a partially realized investment ends at the amount of value that remains unrealized. For investments whose gross value rises above invested capital before realization, this weighting produces the familiar rise and fall NAV profile: value increases early in the holding period and declines as realization shifts weight toward the remaining unrealized value. The bowed transition is motivated by Takahashi and Alexander (2002) and the subsequent empirical calibration of the bow parameter using private market fund data in Shen et al. (2020).

A simple example is shown in Figure A1. In both panels, synthetic NAV initially rises because the gross value path receives most of the weight early in the holding period. As the investment matures, the weight gradually shifts toward the unrealized value path, causing synthetic NAV to peak and then decline. In the full exit case, that decline continues to zero by the terminal date, while in the partial exit case synthetic NAV converges to the remaining unrealized value.

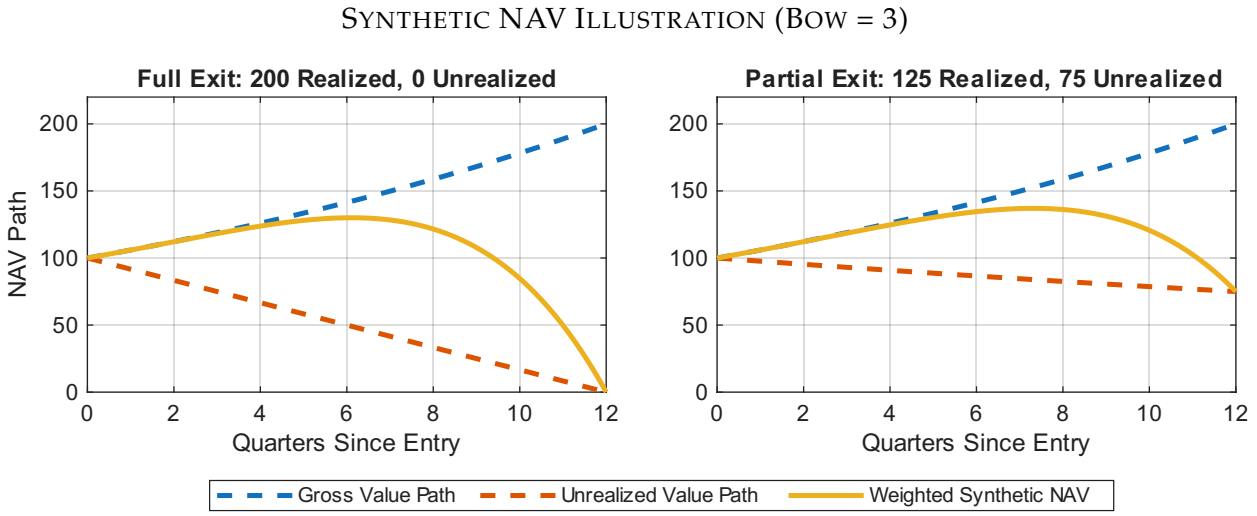


FIGURE A1: This figure shows the synthetic NAV construction for a hypothetical investment with 100 of invested capital and terminal gross value of 200 over a 12 quarter holding period. The dashed lines show the gross value and unrealized value paths, while the solid line shows their weighted combination using a bow parameter of 3. The left panel assumes a full exit with 200 realized and no remaining unrealized value; the right panel assumes a partial exit with 125 realized and 75 remaining unrealized.

For annual exposure analyses, we convert the continuous synthetic path into a time-weighted annual NAV measure. Within each calendar year, the investment is evaluated at the midpoint of the portion of each quarter during which it is active. The synthetic NAV at that midpoint is multiplied by the investment's active days in that quarter as a fraction of the calendar year, and the quarterly contributions are summed. The resulting time-weighted synthetic NAV should therefore be interpreted as an annual capital exposure measure: an investment held for only part of a year contributes proportionally less than an otherwise identical investment held throughout the year, and the measure is therefore distinct from the point-in-time year-end measure described next.

The procedure also calculates a separate year-end synthetic NAV when a point-in-time measure is required. For each calendar year, NAV is evaluated at December 31 when the investment

remains active through year-end; if the observed holding period ends earlier, the terminal date is used instead. If an investment has exited on or before year-end, its year-end NAV is set to zero. No synthetic NAV is projected beyond the observed exit or report date. Downstream analyses can therefore use the time-weighted measure when the object of interest is exposure over the course of a year and the year-end measure when the object of interest is capital remaining at or near the end of the calendar year.

Appendix G. Prospective capital and capital expenditure forecasts

Fund capital. The following describes the methodology used to construct the estimates reported in Table 9. We combine 20 reconstructed deal-level portfolio panels with a broader fund-level fundraising universe to estimate the amount of fund capital attributable to the exposed sectors. Within each reconstruction, investments are classified as either exposed or other, and each general partner's exposed-sector share is estimated from its reconstructed 2023–2025 holdings using NAV, measured as the share of total NAV attributable to the exposed sectors. This share is then matched to the fund-level fundraising data, which provide the capital associated with each general partner's funds.

Using these matched general-partner-level estimates, we next estimate the amount of fund capital attributable to the exposed sectors for each vintage year from 2023 through 2026. Fund capital is measured using reported fund size for closed funds and target fund size for funds that are still raising or in pre-marketing. Within each reconstruction, each general partner's fund capital is multiplied by its estimated 2023–2025 exposed-sector share, and the resulting amounts are pooled across general partners. Because the reconstructed deal-level panels cover only a subset of the general partners represented in the broader fund-level fundraising universe, these pooled estimates are scaled using the corresponding fund-capital coverage adjustment. The procedure is carried out separately for each of the 20 reconstructions, and the resulting estimates are averaged to reduce dependence on any single reconstructed portfolio. The difference between total fund capital and estimated exposed-sector fund capital is classified as capital in the other sectors. The 2026 data are observed through September 30, 2026 at the latest, and the resulting fund-capital estimates are annualized based on the fraction of the calendar year elapsed. Within the exposed total, capital can be allocated across the individual sectors according to each year's active exposed-

sector NAV composition, with the 2026 composition held at its 2025 level; the main text reports the exposed total only, since the sector split of capital that has been raised but not deployed is an assumption rather than an observation. The estimates should be interpreted as an exposure-based measure of fundraising associated with sectors positioned to participate in the physical infrastructure buildout surrounding AI, rather than as a direct measure of capital explicitly raised for AI infrastructure.

Deployable capital and dry powder. The following describes the methodology for creating Table 10. We first calculate cumulative deployed capital through year-end 2023, 2024, and 2025 and decompose it into exposed-sector and other investment. Because cumulative deployed capital includes investments that have subsequently exited, we also calculate the exposed-sector share of time-weighted active synthetic NAV for each year. For each year, we aggregate positive time-weighted synthetic NAV across investments belonging to the eligible funds and divide exposed-sector NAV by total active NAV. To estimate remaining dry powder at year-end 2025, we calculate each fund's remaining deployable capital as 90 percent of reported fund size less cumulative deployed capital through year-end 2025, with negative values set to zero.[24] The 90 percent adjustment allows for a portion of committed capital to be used for management fees and other fund expenses rather than investments. We then allocate each fund's estimated dry powder between the exposed and the other sectors according to the fund's cumulative invested-capital mix, including exited investments, thereby holding the fund's historical investment allocation fixed rather than assuming an additional shift toward the exposed sectors. We measure the difference in exposed-sector shares between projected year-end 2025 dry powder and time-weighted active synthetic NAV during 2025, rather than cumulative deployed capital. We perform the calculation separately across each of the 20 reconstructed fund panels and report the average results.

Illustrative 2026 deployment. The following procedure produces the deployment estimates reported in Table 11 and Table A11. To translate prospective 2026 fund capital into expected 2026 deployment, we use 2023–2025 as a recent historical benchmark. Over these years, deal-level deployment is observed directly, while the reconstructed fund panels allow us to measure beginning-of-

[24]Negative estimated dry-powder values, which may arise from recycling, fund-level borrowing, or differences between reported fund size and cumulative investment measures, are set to zero.

year dry powder and deployable capital associated with funds of that year's vintage. Consistent with the dry-powder methodology, we assume that 90 percent of reported fund size is deployable, with the remaining 10 percent reserved for fees and expenses. Across 2023–2025, funds deploy an average of 11.5 percent of beginning-of-year dry powder and 17.8 percent of new-vintage deployable capital during the year. Applying these rates to the corresponding 2026 capital pools produces the fund-capital-basis deployment estimates reported in Table A11.

Because the deal-level sample covers a narrower universe than the fund-level data, we then translate the fund-capital-basis estimates to a deal-level basis. We calculate a single pooled 2023–2025 deal-to-fund scaling factor as total observed deal-level deployment divided by total deployment implied by the fund-capital methodology. The resulting scaling factor is 0.540. We apply this common factor to all sectors, preserving the relative sector allocations implied by the fund-capital estimates while adjusting for the overall difference between observed deal-level and fund-implied deployment. Applying this factor to the 2026 fund-capital-basis projections produces the deal-level estimates reported in Table 11.

Table A11 reports the corresponding estimates on the fund-capital basis. For 2023–2025, these estimates are constructed by applying historical average deployment rates to beginning-of-year dry powder and new-vintage investable fund capital, rather than by rescaling observed deal-level deployment. The 2026 values are projected deployment amounts on the same basis. Total deployment rises from roughly $94 billion annually over 2023–2025 to $123 billion in 2026, and deployment in the exposed sectors from roughly $35 billion to $64 billion, or from about 37 percent to 52 percent of the total. The historical shares differ from those of Table 11 because the relationship between fund-implied and observed deployment differs between exposed and other sectors. Across 2023–2025, observed deal-level deployment totals roughly $151 billion against $281 billion implied by the fund-level data, or 54 percent, whereas for the exposed sectors the corresponding figures are $45 billion against $104 billion, or 43 percent. For 2026, we apply a common scaling factor of 0.540 across all sectors to obtain the deal-level projections, preserving their relative shares and rankings between the two bases.

TABLE A11: Illustrative 2026 Deployment in Private Real Assets, Fund-Capital Basis

Deployment (USD billions)	**2023**	**2024**	**2025**	**2026**
Total real assets	97.7	87.4	95.4	*123.4*
Exposed sectors	37.5	26.3	39.8	*63.9*
IPP & Renewables	18.3	12.2	16.1	*24.0*
Data Centers	7.3	6.4	11.8	*19.7*
Midstream Energy	8.7	6.2	9.0	*14.1*
Other Exposed Sectors	3.2	1.5	2.9	*6.1*
Exposed share	38%	30%	42%	*52%*

Deployment is reported in USD billions on the fund-capital basis. Historical deployment for 2023–2025 is estimated by applying historical average deployment rates to beginning-of-year dry powder and new-vintage investable fund capital, rather than by rescaling observed deal-level deployment. Other exposed sectors represent the difference between total exposed-sector deployment and the three separately reported subsectors. The 2026 figures (italics) are the projected fund-capital-basis deployment amounts described in the text.

Capital expenditure forecast levels. Table A12 reports the evolution of I/B/E/S consensus capital expenditure forecasts for the same balanced firm sample used in Table 12, by group and fiscal year, from the initial consensus observation to the latest available forecast. To ensure comparability across fiscal years and forecast horizons, the sample is restricted to firms with available initial and latest forecasts, as well as valid February and August forecasts, in every fiscal year from 2024 through 2028. The latest forecast for each firm is selected no later than its fiscal-year end or the reporting cutoff date, whichever occurs first. The resulting firm sample is identical across both tables, and the number of firms in each group is reported. The largest upward revisions are concentrated in the demand complex, for which expected capital expenditure for fiscal 2026 rises from $202 billion to $630 billion and for fiscal 2027 from $313 billion to $930 billion. The exposed real-asset sectors are revised upward in every fiscal year, with cumulative revisions ranging from 9.8 to 39.3 percent, whereas the remaining real-asset sectors show revisions of both signs. Because the interval between the initial and latest observations differs across firms and fiscal years, the cumulative revisions do not represent identical information windows. The standardized six-month revisions of Table 12 therefore remain the primary measure used in the analysis.

TABLE A12: Evolution of Capital Expenditure Forecasts by Group, Fiscal Years 2024–2028

			Initial Consensus		Latest Consensus		Change	
Portfolio	**Fiscal Year**	**Firms**	**Date**	**Expected Capex**	**Cutoff Date**	**Expected Capex**	**$bn**	**Cumulative (%)**
AI Buildout Firms	2024	20	Feb 2022	61.4	Jul 2024	49.9	-11.5	-18.7%
AI Buildout Firms	2025	20	Feb 2023	66.0	Jul 2025	72.1	6.0	9.1%
AI Buildout Firms	2026	20	Feb 2024	62.4	Jul 2026	100.1	37.7	60.5%
AI Buildout Firms	2027	20	Feb 2025	78.4	Aug 2026	139.6	61.2	78.0%
AI Buildout Firms	2028	20	Feb 2026	115.3	Aug 2026	158.3	43.0	37.3%
AI Demand Firms	2024	11	Feb 2022	157.4	Jul 2024	193.2	35.8	22.7%
AI Demand Firms	2025	11	Feb 2023	159.1	Jul 2025	307.7	148.6	93.4%
AI Demand Firms	2026	11	Feb 2024	202.1	Jul 2026	629.8	427.7	211.6%
AI Demand Firms	2027	11	Feb 2025	313.5	Aug 2026	930.3	616.8	196.7%
AI Demand Firms	2028	11	Feb 2026	652.1	Aug 2026	1037.3	385.2	59.1%
All AI Firms	2024	34	Feb 2022	229.9	Jul 2024	262.0	32.1	14.0%
All AI Firms	2025	34	Feb 2023	242.1	Jul 2025	407.0	164.9	68.1%
All AI Firms	2026	34	Feb 2024	284.6	Jul 2026	803.1	518.5	182.2%
All AI Firms	2027	34	Feb 2025	421.0	Aug 2026	1181.2	760.2	180.6%
All AI Firms	2028	34	Feb 2026	857.4	Aug 2026	1321.0	463.7	54.1%
Exposed Real-Asset Sectors	2024	40	Feb 2022	67.3	Jul 2024	93.8	26.5	39.3%
Exposed Real-Asset Sectors	2025	40	Feb 2023	84.3	Jul 2025	112.2	27.9	33.1%
Exposed Real-Asset Sectors	2026	40	Feb 2024	95.1	Jul 2026	130.3	35.3	37.1%
Exposed Real-Asset Sectors	2027	40	Feb 2025	112.7	Aug 2026	146.1	33.4	29.6%
Exposed Real-Asset Sectors	2028	40	Feb 2026	134.3	Aug 2026	147.5	13.1	9.8%
Ex-AI Market	2024	627	Feb 2022	636.0	Jul 2024	666.3	30.2	4.8%
Ex-AI Market	2025	627	Feb 2023	659.8	Jul 2025	697.1	37.3	5.6%
Ex-AI Market	2026	627	Feb 2024	692.4	Jul 2026	743.3	50.9	7.4%
Ex-AI Market	2027	627	Feb 2025	715.9	Aug 2026	787.5	71.6	10.0%
Ex-AI Market	2028	627	Feb 2026	772.8	Aug 2026	802.1	29.4	3.8%
Ex-AI Real Assets	2024	29	Feb 2022	76.4	Jul 2024	88.5	12.1	15.8%
Ex-AI Real Assets	2025	29	Feb 2023	82.9	Jul 2025	88.8	5.9	7.1%
Ex-AI Real Assets	2026	29	Feb 2024	92.7	Jul 2026	89.9	-2.8	-3.1%
Ex-AI Real Assets	2027	29	Feb 2025	94.2	Aug 2026	92.7	-1.5	-1.6%
Ex-AI Real Assets	2028	29	Feb 2026	98.4	Aug 2026	95.9	-2.4	-2.5%

This table reports the evolution of I/B/E/S consensus capital expenditure forecasts for fiscal years 2024 through 2028 using a balanced sample of firms. Within each group, the same firms are included in every fiscal year. Each firm must have valid initial and latest forecasts for all five fiscal years, as well as valid February and August forecasts used to calculate the standardized six-month revisions in Table 12. The resulting firm sample is therefore identical across both tables. The Firms column reports the number of firms in each group, which is constant across fiscal years. Initial forecasts are observed in February two years before the corresponding fiscal year. The latest available consensus forecast for each firm is selected no later than the reporting cutoff date or the firm's fiscal-year end, whichever occurs first. Consequently, the cutoff date need not coincide with every firm's final forecast observation. All capital expenditure amounts and dollar changes are reported in billions of U.S. dollars. Cumulative percentage changes are calculated as the ratio of aggregate latest expected capital expenditure to aggregate initial expected capital expenditure, minus one. Group definitions follow Table 12.